Supply Chain Life Cycle Optimization under Uncertainty

MSc Thesis

written by

Philippe van Mastrigt

under the supervision of **Selin Hülagü MSc**, **Dr. Reinout Heijungs**, and **Dr. Neil Yorke-Smith** and submitted to the Board of Examiners in partial fulfillment of the requirements for the degree of

MSc Embedded Systems

at *Delft University of Technology*.

Date of the public defense:	**Members of the Thesis Committee:**
28 January 2026 @ 14h00	Dr. N. Yorke-Smith
	Prof.dr.ir. A.A.J.F. van den Dobbelsteen
	Dr. M.A. Sharifi Kolarijani

TU Delft VU

Would the certainty really be greater for being checked twenty times?

Ludwig Wittgenstein, *On Certainty*, §77

Abstract

For most organizations, the majority of greenhouse gas emissions are Scope 3 emissions embedded in geographically dispersed supply chains. In such settings, environmental and economic impacts, as well as operating conditions, are uncertain, and decisions are sequential, meaning that early commitments constrain later operational choices. Yet current optimization practices do not integrate a full life cycle perspective in such conditions. In response, this thesis develops a stochastic Closed-Loop Supply Chain Life Cycle Optimization (SCLCO) framework that embeds the algebraic, matrix-based structure of life cycle assessment directly within a multistage stochastic mixed-integer program. Rather than assessing impacts ex post, the model minimizes expected environmental and economic impacts through sequential decisions while representing uncertainty in supplier-dependent life cycle inventory data and economic drivers, as well as stochastic demand and return processes. The framework is illustrated using a university campus circular furniture procurement case study. Applying the stochsatic SCLCO to a framework agreement yields decision-relevant insights for supplier selection under joint environmental–economic objectives. The resulting problem is solved using Sample Average Approximation. Trade-offs are explored via weighted objectives and summarized using a Pareto frontier; robustness is assessed through optimality-gap analysis and value-of-information metrics; and outcomes are examined using contribution, uncertainty, and sensitivity analyses. Results indicate that economic outcomes are driven primarily by demand uncertainty, whereas environmental outcomes are dominated by uncertainty in manufacturing-stage inventory data. Accordingly, this work offers a decision-support framework for decarbonization efforts throughout supply chains.

Acknowledgements

I would like to thank my supervisors, Selin Hülagü, Reinout Heijungs, and Neil Yorke-Smith, for their support throughout this project. From its early, fragile stage through to completion, they encouraged me to explore and refine the research direction, while providing thoughtful feedback and guidance that helped me develop both clarity and confidence.

I also thank my colleagues across the university's faculties and departments for discussions, insights, and a collegial environment that made this work possible. The collaborative spirit and shared commitment to the university's long-term objectives created the conditions under which this research could take shape.

As such, this work is a product of its environment.

List of Abbreviations

Notation	Description	Page List
ADI	average demand interval	24, 57
AIC	Akaike information criterion	60
BEV	battery-electric vehicle	54
BoM	bill of materials	51, 53, 91
CDF	cumulative distribution function	26, 35, 38
CREFM	Campus Real Estate & Facility Management	48
CVaR	conditional value at risk	93
EEV	expected value of the expected value	77
EF	extensive form	23, 64
EPD	environmental product declaration	51, 93
EV	expected value	20
EVPI	expected value of perfect information	10, 11, 20, 39, 40, 41, 64, 76, 77, 78, 87, 89
GHG	greenhouse gas	9, 47, 71, 72, 87, 88
GSA	global sensitivity analysis	16
GWP	global warming potential	91, 92
HEI	higher education institution	24, 47
IPCC	Intergovernmental Panel on Climate Change	
ISO	International Organization for Standardization	13
KDE	kernel density estimator	38, 107
KPI	key performance indicator	51
LB	lower bound	75
LCA	life cycle assessment	9, 11, 13, 14, 15, 16, 17, 18, 27, 50, 62, 64, 87, 92, 94
LCI	life cycle inventory	13, 14, 33, 42, 43, 92, 93
LCIA	life cycle impact assessment	14, 33, 87, 92
LCO	life cycle optimization	17, 18
LP	linear programming	14
MC	Monte Carlo	15, 16, 21, 22, 32, 33, 37, 38, 40, 42, 43, 44, 45, 50, 53, 54, 55
MILP	mixed-integer linear program	18, 20, 23, 40
NAC	non-anticipativity constraint	23, 37, 62
NRMSE	normalised root mean squared error	
PDF	probability distribution function	26
PH	progressive hedging	23, 24, 64, 65
RER	ecoinvent Europe region	
RP	recourse problem	20, 77, 78

Notation	Description	Page List
SAA	sample average approximation	10, 11, 21, 22, 30, 37, 38, 39, 40, 41, 42, 74, 75, 76, 88
SCLCO	Supply Chain Life Cycle Optimization	9, 13, 18, 47, 48, 88, 91, 92
SLA	service level agreement	91
TU Delft	Delft University of Technology	11, 47, 48, 49, 87, 92
UA	uncertainty apportioning	15, 16, 17, 42
UB	upper bound	75
VSS	value of the stochastic solution	10, 11, 20, 39, 40, 64, 77, 78, 87, 89
WS	wait-and-see	20, 77

List of Symbols

Notation	Description
$\mathbf{A}$	Matrix defining the structural input-output relationships of unit processes
$\mathbf{B}$	Matrix defining environmental and economic exchanges per unit of process activity
β	Weibull shape parameter governing the hazard rate of product returns
$\tilde{\mathbf{B}}(x)$	Stochastic intervention matrix associated with supplier x
$\mathbf{e}$	Binary vector indicating which flows must satisfy strict equality constraints
η	Weibull scale parameter governing the characteristic life of products
$\mathbf{f}$	Vector specifying the required final demand for products and services
$\tilde{\mathbf{f}}_\tau$	Stochastic demand vector realized at stage τ
G	Difference between the objective value of a candidate solution and the true optimal value
$\mathbf{g}$	Calculated life cycle inventory flows
$\tilde{\mathbf{g}}_\tau$	Stochastic inventory vector realized at stage τ
$h(x;\lambda)$	Function evaluating the expected aggregate impact for a given supplier and weight
$\tilde{H}(x;\lambda)$	Random variable representing the total weighted impact
i	Index identifying a specific process or flow
$\mathbf{K}$	Leslie matrix operator handling the inter-temporal carry-over of inventory
k	Index identifying a product category
λ	Scalar weight controlling the trade-off between environmental and economic objectives ($\lambda \in [0,1]$)
N_c	Number of capacity constraints
N_e	Total number of interventions
N_f	Total number of flows
N_i	Total number of impact categories
N_i^{ecn}	Number of economic impact categories
N_i^{env}	Number of environmental impact categories
N_p	Total number of unit processes
Ω	Set of all scenario paths in a scenario tree
ω	Index identifying a specific scenario path
$\Omega^{(r)}$	Set of scenario paths belonging to tree replication r
$\mathbf{P}$	Combined matrix containing both technology coefficients and intervention data
p_ω	Probability mass associated with scenario path ω
π_r	Probability of selecting scenario tree replication r
$\mathbf{Q}$	Characterization matrix mapping elementary interventions to life cycle impact categories
R	Total number of independent scenario tree replications
r	Index identifying a scenario tree replication

Notation	Description
$\mathbf{s}$	Vector of decision variables representing process activity levels
T	Total number of decision stages in the planning horizon
τ	Index denoting the time stage, $\tau \in \{1, \ldots, T\}$
$\mathbf{u}(\mathrm{x})$	Vector defining upper bounds on output flows for supplier x
$\mathbf{W}(\lambda)$	Operator that aggregates multiple impact categories into a weighted objective
$\mathcal{X}$	Set of available suppliers for selection
x	First-stage decision variable representing the choice of supplier
ξ	Random vector capturing all stochastic parameters (demand, returns, impacts)
z_{eev}	Expected result of using the Expected Value solution in the stochastic environment
z_{ev}	Optimal objective value of the Expected Value problem (using mean parameters)
z_{rp}	Optimal objective value of the Recourse Problem (stochastic solution)
z_{ws}	Expected objective value of the Wait-and-See solution (perfect information)

Contents

Chapter 1

Introduction

For decades, climate change researchers have consistently highlighted the severe consequences of global warming, underscoring the urgent need for comprehensive and effective mitigation efforts (Calvin et al., 2023). Accordingly, under the 2015 Paris Agreement, many countries committed to limiting the global average temperature increase to 1.5 °C above pre-industrial levels (Paris Agreement, 2015). In turn, an increasing number of public and private organizations have adopted long-term decarbonization targets and have begun systematically monitoring their greenhouse gas (GHG) emissions, driven both by regulatory disclosure requirements and by strategic planning needs.

Despite these efforts, translating long-term climate ambitions into actionable and measurable short- to medium-term outcomes remains a significant challenge. In particular, for most organizations, a substantial share of emissions lies beyond direct operational control and is embedded within complex and geographically dispersed value chains, commonly referred to as Scope 3 emissions. Reducing these indirect emissions is difficult because data are limited and heterogeneous, reporting standards remain inconsistent, supplier engagement is challenging, mitigation options may be economically constrained, and monitoring and forecasting are resource-intensive (Carrillo Pineda et al., 2023). These challenges highlight that supply chain decisions are subject to substantial uncertainty in both environmental and economic outcomes, especially in sequential settings where early commitments restrict later operational choices.

Although the literature on mitigation and decarbonization has increased considerably, the integration of environmental and economic criteria into supply chain decision-making in the above conditions remains limited. Life cycle assessment (LCA) provides a comprehensive framework for quantifying impacts across product life cycles; nevertheless, it is most often applied ex post to fixed system configurations and is rarely incorporated endogenously into decision models. By contrast, supply chain optimization models support strategic and operational planning, yet they typically cannot represent complete life cycle inventories, and environmental performance is commonly treated deterministically, with uncertainty modeling often confined to exogenous drivers such as demand or prices. Consequently, existing approaches provide limited support for supply chain decisions when both life cycle inventories and operational conditions are uncertain. Therefore, given the severity of climate change consequences, methods are needed that embed a life cycle perspective within supply chain decision models and jointly optimize environmental and economic performance under uncertainty.

1.1 Research Objectives

In response, this thesis develops and evaluates a decision-support methodology that embeds a life cycle perspective within multistage supply chain planning, jointly optimizing environmental and economic performance under uncertainty. Specifically, it proposes a stochastic Closed-Loop Supply Chain Life Cycle Optimization (SCLCO) framework, a multistage stochastic optimization model in which the algebraic structure of LCA is em-

bedded within the decision model itself. Rather than evaluating supply chain impacts after decisions are fixed, the framework treats life cycle impacts as endogenous model components that shape optimal supply chain configurations and recourse decisions over time.

The decision context studied in this thesis involves interacting strategic and tactical choices. A strategic *here-and-now* decision is made at the beginning of the planning horizon, such as selecting a supplier under a long-term framework agreement. This decision is followed by a sequence of periodic *recourse* decisions that adapt to realized uncertainty and determine optimal actions, including choices related to manufacturing technologies, (reverse) logistics, and end-of-life circular processing. Uncertainty is represented in two forms: (i) supplier-dependent environmental and economic variables specified ex ante, including life cycle inventory and cost drivers, and (ii) exogenous interstage stochastic processes, such as demand and return flows, that evolve over time. The framework is demonstrated in an organizational procurement context through a circular campus furniture case study consistent with this sequential structure.

The resulting stochastic optimization problem is solved using sample average approximation (SAA), yielding a deterministic equivalent mixed-integer formulation that supports scenario-based evaluation of expected outcomes. Trade-offs between environmental and economic performance are examined using weighted objective formulations. Robustness is assessed through gap analysis, contribution analysis, and uncertainty apportioning. Decision-value metrics, including the value of the stochastic solution (VSS) and the expected value of perfect information (EVPI), are used to evaluate the benefit of explicitly modeling uncertainty relative to deterministic planning approaches.

Crucially, uncertainty is not treated merely as variability around fixed outcomes. In multistage supply chain settings, uncertainty affects which decisions are optimal in the first place because early commitments constrain future recourse decisions. Ignoring uncertainty can therefore yield policies that appear effective under assumed average conditions but prove fragile under plausible future realizations. By explicitly modeling uncertainty, the proposed framework can identify hedging policies, quantify the value of information, and provide sensitivity-based guidance on which sources of uncertainty most influence decisions and outcomes.

1.2 Research Questions

The study is guided by three general research questions:

Research Question 1: *How can uncertainty be methodologically integrated into a supply chain life cycle optimization framework that jointly optimizes environmental and economic impacts?*

Research Question 2: *How can the proposed framework be applied to a real-world organizational context characterized by uncertain operational constraints, supplier performance, and demand flows?*

Research Question 3: *What trade-offs and policy insights emerge from the framework, and to what extent are these outcomes robust when assessed through systematic uncertainty analysis?*

1.3 Research Approach

The research approach is structured directly around the three research questions. To address the methodological question (1), a multistage stochastic closed-loop supply chain life cycle optimization framework is developed. This involves three key modeling steps: (i) constructing life cycle inventories and associated impact models in matrix form; (ii) characterizing relevant sources of uncertainty, including probability distributions of input

variables and exogenous stochastic processes; and (iii) formulating a stochastic mixed-integer optimization model that embeds the algebraic LCA structure endogenously. The resulting model is solved using SAA, and solution quality is evaluated using statistical bounding techniques.

To address the empirical question (2), the framework is applied to a concrete organizational procurement context of Delft University of Technology (TU Delft). This case study demonstrates how strategic supplier selection and tactical recourse decisions, such as production scaling, reverse logistics, and end-of-life processing, can be optimized jointly under realistic constraints. The application serves as an illustration for the methodology's capacity of handling complex, demand patterns and circular flows.

Finally, to address the third research question (3), the optimal policies generated by the framework are subjected to post-optimality analysis. Decision insights are derived by assessing policy robustness via optimality-gap analysis, quantifying the value of explicitly modeling uncertainty using metrics such as the VSS and the EVPI, and identifying contributions and influential uncertainty sources through variance decomposition.

1.4 Contributions

Building on the research questions and outlined approach, its contributions are threefold.

First, it develops a unified methodological framework that integrates the algebraic structure of LCA directly into a multistage stochastic closed-loop supply chain optimization model. By embedding life cycle inventory and impact assessment within the decision model, environmental and economic impacts become endogenous to both strategic design choices and intertemporal recourse decisions, rather than being evaluated ex post for fixed configurations. This integration enables the joint treatment of uncertainty in supplier-dependent life cycle variables (inventory uncertainty) and uncertainty in stochastic exogenous drivers (stochastic process uncertainty), within a single coherent optimization framework.

Second, the thesis demonstrates the applicability of the proposed framework in a real-world organizational procurement setting aimed at reducing environmental and economic impacts. The case study links a long-term *here-and-now* commitment (supplier selection) to multi-period recourse decisions in manufacturing, logistics, and end-of-life processing under uncertain demand and returns. By embedding environmental and economic impacts together with circularity options and operational constraints (e.g., take-back, reuse and recycling pathways, and capacity or policy limits), the application provides an illustration on how to quantitatively address supply-chain decarbonization efforts.

Third, it contributes an evaluation and decision-support perspective for interpreting stochastic life cycle optimization results under uncertainty. The thesis combines optimality gap analysis with contribution analysis, uncertainty apportioning, and value-of-information metrics, including the VSS and the EVPI, to assess robustness and decision relevance. This assessment clarifies how to evaluate the framework, identifying which inputs and stochastic drivers most influence performance trade-offs, and provides a structured basis for prioritizing data collection and risk management efforts in life-cycle-informed supply chain planning.

1.5 Structure of the Thesis

The remainder of this thesis is organized as follows. Chapter 2 reviews the theoretical background, covering life cycle assessment, matrix-based computational structures, uncertainty and sensitivity analysis, and supply chain optimization under uncertainty. Chapter 3 presents the methodological framework, formulating the stochastic closed-loop supply chain model and describing the computational approach based on SAA. Chapter 4 introduces the empirical case study, detailing the problem specification, data sources, life cycle modeling choices, and the generation of stochastic demand and return scenarios. Chapter 5 presents the optimization outcomes, including trade-off analyses, optimality

gaps, and uncertainty quantification. Finally, Chapter 6 interprets the findings, discusses limitations and practical implications, and concludes with answers to the research questions and recommendations for future research.

Chapter 2

Theoretical Background

This chapter provides a review of the theoretical and computational basis required to develop and evaluate the stochastic SCLCO framework. It bridges the fields of environmental life cycle assessment, sustainable supply chain management, and stochastic optimization. Section 2.1 introduces LCA and its matrix-based formalization. Section 2.2 discusses uncertainty quantification and sensitivity analysis in the LCA context. Section 2.3 reviews existing literature on sustainable supply chain optimization, highlighting current limitations, such as uncertainty treatment. Section 2.4 details the principles of multistage stochastic programming and value-of-information metrics. Finally, Section 2.5 covers demand forecasting and reliability modeling, which serve as the basis for generating stochastic scenarios for the empirical case study.

2.1 Life Cycle Assessment

LCA is a method for evaluating the environmental impacts of products, processes, and services across their entire life cycle. By accounting for impacts from raw material extraction to end-of-life, LCA provides a basis for understanding environmental burdens and supporting informed decision-making. An accessible, general, and recent introduction on LCA and its life cycle perspective is provided by Kaynak et al. (2025).

Its practice helps organizations to identify improvement opportunities and develop more resilient production systems by assessing the full supply chain and relevant impact categories. Its broad scope distinguishes it from methods that focus on single stages or indicators, and it is applied across a wide range of industrial and environmental contexts, from single-product studies to broader organizational or territorial assessments.

LCA is commonly understood through the International Organization for Standardization (ISO) 14040 standard (ISO 14040:2006(E), 2006), which provides guidelines for consistency and transparency. These guidelines call for defining the study's purpose, system boundaries, and functional unit, followed by compiling a life cycle inventory (LCI) that quantifies material and energy flows. These flows are then linked to environmental impact categories through established modeling approaches. The results are interpreted to identify significant contributors, evaluate uncertainties, and draw conclusions that can inform further analysis or decision-making.

Even within these standards, the many unavoidable methodological choices (e.g., system boundaries, functional unit, allocation procedures, impact assessment methods, software) introduce deviations that inevitably influence results and interpretation. In the following subsection, the computational structure LCA is addressed, necessary to understand the methodological approach outlined in subsequent chapters.

2.1.1 Computational Structure

In this study, the assessment is carried out on the basis of an explicit algebraic structure, as outlined in other works (Heijungs et al., 2013; Heijungs & Suh, 2002; Suh et al., 2010).

An LCA system can be defined in terms of (unit) processes and inventory quantities, $\mathbf{P} \in \mathbb{R}^{(N_f+N_e)\times N_p}$ and $\mathbf{q} \in \mathbb{R}^{N_f+N_e}$ respectively, for which the following notation is used:

$$\mathbf{P} = \left[\frac{\mathbf{A}}{\mathbf{B}}\right], \qquad \mathbf{q} = \left[\frac{\mathbf{f}}{\mathbf{g}}\right]. \tag{2.1}$$

Here, N_p, N_f, and N_e denote the total number of unit processes, flows, and environmental interventions, respectively; $\mathbf{A} \in \mathbb{R}^{N_f \times N_p}$ denotes intra-system exchanges, also referred to as the technology matrix or technosphere; $\mathbf{B} \in \mathbb{R}^{N_e \times N_p}$ denotes environmental interventions, also known as the intervention or biosphere matrix; $\mathbf{f} \in \mathbb{R}^{N_f}$ represents the functional unit, reference demand, or exogenous input; and $\mathbf{g} \in \mathbb{R}^{N_e}$ represents the LCI.

Under linearity, colloquially meaning that doubling the input doubles the output,

$$\mathbf{s} = \mathbf{A}^{-1}\mathbf{f}, \tag{2.2}$$

$$\mathbf{g} = \mathbf{B}\mathbf{s} = \mathbf{B}\mathbf{A}^{-1}\mathbf{f}, \tag{2.3}$$

where $\mathbf{s} \in \mathbb{R}^{N_p}$ denotes the activity levels, also referred to as the scaling vector. Unit consistency follows directly from the dimensional structure of $\mathbf{A}$ and $\mathbf{B}$.

The expressions above define the inventory matrix but do not yet characterize the resulting impacts. For this purpose, life cycle impact assessment (LCIA) is applied, in which the inventory is translated into impact categories.

$$\mathbf{h} = \mathbf{Q}\,\mathbf{B}\,\mathbf{A}^{-1}\mathbf{f}, \tag{2.4}$$

where N_i denotes the total number of impact categories, $\mathbf{h} \in \mathbb{R}^{N_i}$ is the vector of impact scores, and $\mathbf{Q} \in \mathbb{R}^{N_i \times N_e}$ maps inventory flows to N_i impact indicators (e.g., IPCC 2021 GWP_{100} for climate change), with classification and normalization applied as required.

This above LCA formulation raises concerns such as cut-off or multifunctionality, or put in algebraic terms, raises issues of invertibility when $\mathbf{A}$ is rectangular. Its classical resolutions can be found in the literature (Heijungs & Suh, 2002).

In this study, linear programming (LP) replaces the problem of defining the system as an inversion problem by instead formulating it as an optimization problem, see Heijungs (2024a, Sec. 8.6.3). The solution selects scaling vectors such that they satisfy the balance equation

$$\mathbf{A}\mathbf{s} = \mathbf{f},$$

while minimizing (or maximizing) a specified objective function (Budzinski et al., 2019). A generic LP formulation for minimizing impacts is then given by:

$$\begin{aligned} \min_{\mathbf{s}} \quad & \mathbf{Q}\,\mathbf{B}\,\mathbf{s} \\ \text{s.t.} \quad & \mathbf{A}\,\mathbf{s} = \mathbf{f}, \\ & \mathbf{s} \geq 0. \end{aligned} \tag{2.5}$$

Here, the objective $\mathbf{Q}\,\mathbf{B}\,\mathbf{s}$ represents total impacts associated with activity levels $\mathbf{s}$, subject to technological balance constraints and non-negativity of activities. This is the starting point of the computational structure.

2.2 Uncertainty and Sensitivity Analysis in LCA

Within LCA, uncertainty and sensitivity analysis are central to interpreting results and evaluating their reliability. Uncertainty arises from imperfect knowledge and inherent variability in the underlying data, encompassing measurement error, model simplifications, missing information, and genuine fluctuations in technological or environmental systems (Bamber et al., 2019; Huijbregts et al., 2003). Accordingly, numerical inputs in an LCA model should not be treated as fixed, true quantities but as nominal values representing a range of plausible outcomes. Treating them as exact neglects the variability

and ambiguity inherent in both data and modeling assumptions. Uncertainty analysis is therefore necessary for characterizing the dispersion of possible outcomes and quantifies the extent to which this variability shapes the resulting life cycle impacts (Heijungs, 2024a; Wei et al., 2014).

Contemporary literature emphasizes that uncertainty in LCA is multidimensional, extending beyond input parameter variability to include scenario (or choice) uncertainty arising from methodological decisions, such as system boundaries, allocation methods, functional units, and technology pathways, as well as model uncertainty stemming from differences in characterization models and impact assessment methods (Barahmand & Eikeland, 2022; Beltrán et al., 2015; Zelm & Huijbregts, 2013). These forms of uncertainty often dominate overall variability yet remain insufficiently quantified in practice (Bamber et al., 2019).

Sensitivity analysis complements uncertainty analysis by examining how variations in inputs translate into variations in outputs. It identifies which data points, modeling assumptions, or methodological choices exert the greatest influence on results. While many studies still rely on local, one-at-a-time approaches, such methods overlook correlations and interaction effects within the life cycle model (Wei et al., 2014). More comprehensive global sensitivity methods, such as variance-based techniques like Sobol and Shapley, allow the attribution of output variance to individual inputs or groups of inputs, offering a more accurate representation of influential parameters.

Together, uncertainty and sensitivity analyses provide a basis for assessing robustness and identifying key issues. Inputs that display both high uncertainty and to which the output is sensitive are priority targets for improved data collection or methodological refinement. This process, termed uncertainty apportioning (UA) (Heijungs, 2024a), supports transparent interpretation by allocating the contribution of different uncertainty sources to the overall result variance (Barahmand & Eikeland, 2022). In doing so, it directly enhances the credibility, defensibility, and decision-support value of LCA outcomes, especially in contexts where large variations between studies are common and misinterpreted as methodological unreliability (Ross & Cheah, 2018; Williams et al., 2009).

2.2.1 Uncertainty Quantification and Propagation

Within the algebraic formulation, uncertainty enters the model through the technology matrix $\mathbf{A}$, the intervention matrix $\mathbf{B}$, and the characterization matrix $\mathbf{Q}$. In the LP formulation, uncertainty in any of these matrices changes the optimization problem itself, altering the feasible set or the objective, and therefore leads to different optimal scaling vectors $\mathbf{s}$, different inventories $\mathbf{g}$, and ultimately in different impact results $\mathbf{h}$. Uncertainty analsysis, thus, answers what is the spread, range, or distribution of the outcomes. But in order to come to an answer, computational and quantification methods are required.

The propagation of uncertainty can be assessed in roughly two ways: analytically or by simulation. Analytical approaches, such as first-order Taylor expansion or matrix-based propagation, are efficient but rely on linearity, independence, and small-uncertainty assumptions that are often violated in practice where uncertainties are large (Imbeault-Tétreault et al., 2013). Simulation approaches, most commonly Monte Carlo (MC) sampling, can deal in a non-linear fashion with interactions among input variables, yet require explicit probability distributions and their distribution parameters for all uncertain inputs; which can become computationally demanding in large background systems (Hung & Ma, 2008).

For this study, specifying coherent probabilistic representations of inventory input variables is a prerequisite step for incorporating uncertainty into the methodology. A widely used framework is the pedigree matrix (Weidema & Wesnæs, 1996), providing a means to document data quality, a basis for estimating uncertainty in these input variables when statistical data are missing, and a way to compare the quality of different data sources. It rates data quality along several dimensions, including reliability, completeness, temporal correlation, geographical correlation, and technological correlation. These qualitative data-quality indicators are then translated into a geometric standard deviation for

a lognormal distribution, that is, into distribution parameters describing the uncertainty assigned to input variables (Müller et al., 2014), and is widely used in LCA practices and process databases like ecoinvent (Ciroth et al., 2013). Other types of probability distributions are also employed and present in these databases, such as uniform, triangular, and normal distributions (Müller et al., 2014; Qin & Suh, 2016).

Despite its widespread use, the use of pedigree scores in uncertainty quantification suffers from limitations: it reflects only uncertainty in the input variables (often referred to as 'parameter uncertainty') and not methodological, scenario, or structural uncertainty (Bamber et al., 2019); it relies on subjective and inconsistently applied scoring; its mappings from qualitative scores to geometric standard deviations lack empirical justification (Müller et al., 2014); and it combines 'regular' with pedigree-derived variances through ad-hoc additive rules with no firm statistical basis (Ciroth et al., 2013). Moreover, pedigree factors are frequently misinterpreted as describing process variability, that is, as true distribution parameters, rather than as meta-parameters that reflect data quality and epistemic uncertainty (Barahmand & Eikeland, 2022), leading to inconsistent and sometimes contradictory lognormal parameterizations across LCA software and databases, including unit-dependent distortions and mismatches between specified and realized variance (Heijungs, 2023, 2024a, 2024b), all of which undermine the coherence and statistical interpretability of the resulting uncertainty estimates.

From this perspective, uncertainty analysis relies on assigning coherent probability distributions and propagating them through the model; while the pedigree matrix remains a practical and widely used tool for this purpose, its limitations require careful application and transparent communication of assumptions. In this study, the ecoinvent pedigree-based uncertainties are used as a structured prior over inventory inputs for the stochastic optimization, recognizing that these distributions most often capture epistemic data-quality uncertainty rather than empirical process variability; the resulting solutions should therefore be interpreted conditionally on this uncertainty model, and not as fully robust to all real-world uncertainties.

2.2.2 Global Sensitivity Analysis

Sensitivity analysis complements uncertainty analysis by identifying which inputs or modeling choices drive variability in LCA results. It helps clarify which inputs strongly influence results (even if they are not especially uncertain), illuminates dominant mechanisms, detects non-linearities and interactions, and supports transparent communication of what drives conclusions (Groen et al., 2016; Heijungs, 2024a). Global sensitivity analysis (GSA) explores the full uncertainty space by varying all inputs simultaneously and therefore captures effects that local one-at-a-time methods cannot detect (Wei et al., 2014).

A prominent family of GSA methods is variance-based sensitivity analysis, which attributes portions of the output variance to inputs or groups of inputs. Under standard independence assumptions, Sobol' indices decompose variance into main and interaction contributions (Groen et al., 2016). For dependent or correlated inputs, *Shapley effects* provide an alternative variance allocation principle that remains well-defined in the presence of dependence and strong interactions, albeit at higher computational cost (Iooss & Lemaître, 2015; Kim et al., 2025). In addition to variance-based indices, recent high-dimensional LCA studies also use surrogate models together with SHAP (Shapley additive explanations) to obtain scalable sensitivity rankings rather than variance shares (Kim et al., 2025).

2.2.3 Uncertainty Apportioning

UA links uncertainty quantification with sensitivity concepts by allocating a measure of output dispersion (most commonly variance) to uncertainty sources (Heijungs, 2024a). Practically, it combines uncertainty propagation (often via MC simulation) with a variance allocation rule to determine which uncertainties are responsible for the variation observed in model outcomes (Hung & Ma, 2008). This clarifies which uncertainties matter most

for robustness and where improved data or modeling effort would most effectively reduce overall uncertainty (Bamber et al., 2019; Cucurachi et al., 2015).

Many LCA models contain multiple types of uncertainty that are naturally structured in layers, for example, discrete modeling choices (e.g., system modeling options or scenarios), followed by uncertain parameters within the chosen structure (e.g., inventory and characterization parameters). Heijungs (2024a) suggests that a useful way to connect uncertainty analysis with sensitivity reasoning is UA: allocating a measure of output dispersion (often variance) to uncertainty sources or groups of inputs. A central theoretical tool is the law of total variance, which decomposes output variance into a part that remains after conditioning on some information and a part explained by that information:

$$\mathrm{Var}(Y) = \mathbb{E}[\mathrm{Var}(Y \mid X)] + \mathrm{Var}(\mathbb{E}[Y \mid X])\,. \tag{2.6}$$

This identity admits a natural interpretation for apportioning. The term $\mathrm{Var}(\mathbb{E}[Y \mid X])$ measures how much variance is attributable to the uncertainty in X (because it quantifies how the conditional mean changes as X varies), while $\mathbb{E}[\mathrm{Var}(Y \mid X)]$ represents the average residual variance that would remain if X were fixed.

When uncertainty sources are hierarchical or grouped, the same principle can be applied iteratively by conditioning stepwise. Let R denote a coarse (outer) uncertainty source (e.g., a modeling choice) and S a finer (inner) uncertainty source conditional on R (e.g., a scenario or sub-choice), with remaining uncertainty represented implicitly by other inputs. Iterating the law of total variance yields a multi-level variance decomposition:

$$\mathrm{Var}(Y) = \underbrace{\mathrm{Var}(\mathbb{E}[Y \mid R])}_{\text{variance attributable to } R} + \underbrace{\mathbb{E}[\mathrm{Var}(\mathbb{E}[Y \mid R, S] \mid R)]}_{\text{additional variance attributable to } S \text{ given } R} + \underbrace{\mathbb{E}[\mathrm{Var}(Y \mid R, S)]}_{\text{residual variance from remaining sources}}\,.$$

This decomposition is purely probabilistic: it does not require linearity, and it applies equally to discrete and continuous uncertainty sources. It provides variance shares at the level of structured uncertainty sources or groups, which is precisely the objective of UA in the sense emphasized by Heijungs (2024a).

At the level of grouped or hierarchical sources, the decomposition above is fully general; additional assumptions only arise when one further allocates the residual term to individual parameters (e.g., assuming independence to avoid covariance terms), or when using variance-based sensitivity indices whose interpretation depends on the dependence structure among inputs.

2.3 Sustainable Supply Chain Optimization

While LCA provides the tools for quantifying environmental impacts, its integration into operational decision-making requires an optimization perspective. This section describes the existing context of embedding life-cycles in supply chain network design and operational planning models, transforming environmental assessment from a retrospective reporting tool into a proactive driver of supply chain configuration.

Integrating LCA into an optimization framework transforms environmental assessment from an *ex post* accounting step into an active component of decision-making. This approach falls under the broader umbrella of life cycle optimization (LCO), a field that applies mathematical programming to minimize the life cycle impacts of products, processes, or systems. When applied to supply chain management, LCO extends the traditional boundaries of network design by embedding environmental criteria directly into the objective functions and constraints that shape the supply chain configuration. In algebraic terms, the typical LCA formulation, where the technology matrix $\mathbf{A}$, intervention matrix $\mathbf{B}$, and characterization matrix $\mathbf{Q}$ quantify activity levels to impact scores, is embedded inside a mathematical program. This allows for the simultaneous determination of facility locations, capacities, technology choices, and material flows across the network (Guillén-Gosálbez & Grossmann, 2008; You et al., 2011).

The motivation for such integration in the literature is that supply chain decisions have system-wide and long-term consequences that cannot be optimized in isolation. Research

emphasizes that supply chain design is inherently multi-dimensional, involving trade-offs between environmental burdens and economic performance that often conflict (Barbosa-Póvoa et al., 2017; Brandenburg et al., 2014). Moreover, life cycle consequences arise from interactions across extraction, production, transport, use, and end-of-life processes. These interactions are particularly complex in closed-loop supply chains, where circular flows such as reuse, remanufacturing, and recycling create feedback loops between downstream returns and upstream production (Hülagü et al., 2025; Lozano-Oviedo et al., 2024).

Methodologically, sustainable supply chain optimization is typically formulated as a mixed-integer linear program (MILP) in which discrete variables represent structural choices (e.g., supplier selection, technology choice) and continuous variables represent material and product flows (Guillén-Gosálbez & Grossmann, 2008; Yue et al., 2013). Multi-objective formulations are common, with trade-offs explored via Pareto frontiers or scalarization approaches such as weighted sums (Brandenburg et al., 2014; Jayarathna et al., 2021). However, a persistent limitation identified in this domain is the treatment of uncertainty.

Most existing LCO models in the supply chain literature are deterministic, assuming fixed values for demand, costs, and impact factors. This assumption renders models brittle: optimal solutions for nominal parameters may perform poorly or become infeasible under real-world variability. The literature addresses this primarily through two avenues: *robust optimization*, which seeks solutions feasible under worst-case realizations, and *stochastic programming*, which optimizes expected performance over probability distributions (Bairamzadeh et al., 2015; Mota et al., 2017; Ruiz-Femenia et al., 2013). While robust approaches are computationally attractive, they can be overly conservative. Stochastic programming offers a more flexible framework for managing risk, particularly in multi-stage settings where recourse decisions can adapt to unfolding uncertainty (Gao & You, 2017; Shapiro, 2007).

Despite these advances, significant gaps remain. First, uncertainty in LCA variables (the $\mathbf{B}$ and $\mathbf{Q}$ matrices) is rarely modeled jointly with operational uncertainty (e.g., demand) (Bamber et al., 2019; Hülagü et al., 2025), leading to an underestimation of environmental risk. Second, data quality and methodological mismatches between LCA databases and operational decision models often limit the interpretability of results (Kumar & Kumar, 2024). Finally, the computational complexity of combining closed-loop dynamics with uncertainty often forces researchers to rely on simplified static models rather than full multi-stage formulations.

Building on this literature, the SCLCO model of Hülagü et al. (2025) represents a significant step by expressing a multi-period, closed-loop supply chain directly using LCA data structures. The present work extends this line of research by retaining its explicit matrix-based LCA representation while integrating a fully stochastic treatment of impact and demand uncertainty. In doing so, the optimization not only provides impact-efficient supply chain configurations but also yields insight into how uncertainty in life cycle data propagates to optimal solutions and which processes and decisions most influence the impact outcomes.

2.4 Stochastic Programming and Optimization

Uncertainty in data and future conditions makes supply chain optimization difficult. This section presents stochastic programming, emphasizing the extension of two-stage to multistage recourse models for sequential decisions, and evaluation metrics that measure the value of representing uncertainty.

Stochastic programming provides a mathematical framework for optimization under uncertainty in which some model parameters are represented as random variables. In its risk-neutral form, uncertainty enters through an expectation operator: rather than optimizing against a single nominal parameter realization, one optimizes the expected objective (and/or enforces feasibility requirements) with respect to a specified probability model (Shapiro, 2007). This perspective is particularly well-suited to sequential decision

problems in which uncertainty is progressively revealed over time, and decisions must remain implementable given the information available when they are made.

2.4.1 Two-Stage Stochastic Programming

A common setup is the two-stage stochastic program with recourse. It distinguishes between a first-stage decision made before uncertainty is observed (the "here-and-now" decision) and a second-stage decision made after a realization of the uncertain data becomes available (the "recourse" decision) (Shapiro, 2007).

Let $x \in \mathbb{R}^{n_1}$ denote the first-stage decision, ξ a random vector taking values in Ξ, and $y \in \mathbb{R}^{n_2}$ the recourse decision. For a realized ξ, the second-stage (recourse) problem takes the generic form

$$F(x, \xi) := \min_{y \in G(x,\xi)} g(x, y, \xi), \tag{2.7}$$

where $G(x, \xi)$ is a (possibly ξ-dependent) feasible set and g is the second-stage cost.

At the first stage, one chooses x to optimize the expected recourse value:

$$\min_{x \in X} \; f(x) := \mathbb{E}[F(x, \xi)]\,, \tag{2.8}$$

where X encodes deterministic first-stage constraints.

A common special case is the two-stage stochastic linear program with linear recourse. In its standard cost plus expected recourse representation,

$$\min_{x \in X} \; c^\top x + \mathbb{E}[Q(x, \xi)]\,, \tag{2.9}$$

where $Q(x, \xi)$ is the optimal value of the (linear) second-stage program parameterized by x and the realization of ξ.

When the distribution of ξ is approximated by a finite set of scenarios $\{\xi^k\}_{k=1}^K$ with probabilities $\{p_k\}_{k=1}^K$, the expectation becomes a finite sum and the problem can be expressed as a deterministic equivalent, typically introducing one copy of the recourse variables per scenario (Shapiro, 2007).

2.4.2 Multistage Stochastic Programming

Two-stage models assume uncertainty is revealed in a single future block. In many real systems, uncertainty unfolds gradually, and decisions are made repeatedly as new information arrives. Multistage stochastic programming generalizes the two-stage structure to T stages by modeling both the random data and the decisions as stochastic processes (Dupacová et al., 2000; Shapiro, 2011).

Following Dupacová et al. (2000), consider a stochastic data process $\omega = (\omega_1, \ldots, \omega_T)$ whose realizations are data trajectories, and a decision process $x = (x_1, \ldots, x_T)$ that may depend on observed history. Let $\mathcal{F}_t$ denote the σ-field generated by observations up to (and including) stage t. Equivalently, one can denote the observed history by $\xi_{[t]} := (\xi_1, \ldots, \xi_t)$ (notation varies by source), capturing all information revealed up to stage t (Dupacová et al., 2000; Shapiro, 2011).

A widely used multistage linear stochastic programming form can be written as a nested sequence of stage problems under expectations. For example, Shapiro (2011) presents a multistage linear model in which stage-wise costs and constraints depend on a stochastic data process and decisions are made sequentially:

$$\begin{aligned} &\min\; c_1^\top x_1 + \mathbb{E}\Big[\min\; c_2^\top x_2 + \mathbb{E}\big[\cdots + \mathbb{E}[\min\; c_T^\top x_T]\big]\Big] \\ &\text{s.t. (stage-wise linking constraints coupling } x_{t-1} \text{ and } x_t \text{ for } t = 2, \ldots, T), \end{aligned} \tag{2.10}$$

where the random coefficients form a history-dependent process and the history up to stage t is explicitly tracked (e.g., as $n_{[t]}$ in Shapiro, 2011). This nested form makes explicit that each stage optimizes given what has been observed so far, while performance is evaluated in expectation over future uncertainty.

2.4.3 Value of Information

Beyond providing an implementable decision policy, stochastic programming enables quantification of how much uncertainty and information matter in a given application. Two standard metrics are the EVPI and the VSS, which compare the stochastic program to idealized or deterministic benchmarks (Birge & Louveaux, 2011).

Consider a minimization setting with random data ξ and loss $h(x, \xi)$. Let z_{rp} denote the optimal objective value of the recourse problem (RP), the stochastic program, and let

$$z_{\text{ws}} := \mathbb{E}_{\xi}\left[\min_{x} h(x, \xi)\right]$$

denote the wait-and-see (WS) value, i.e., the expected performance if perfect foresight were available. The EVPI is defined as

$$\text{EVPI} := z_{\text{rp}} - z_{\text{ws}} \geq 0,$$

and measures the opportunity cost of uncertainty (Birge & Louveaux, 2011).

A second benchmark replaces random quantities with their expectations. Let x_{ev} solve the expected value (EV) problem obtained by substituting ξ by $\mathbb{E}[\xi]$, and define the expected result of the EV solution as

$$z_{\text{eev}} := \mathbb{E}_{\xi}[h(x_{\text{ev}}, \xi)]\,.$$

The VSS is

$$\text{VSS} := z_{\text{eev}} - z_{\text{rp}} \geq 0,$$

quantifying the improvement of the stochastic solution over the deterministic EV approximation (Birge & Louveaux, 2011). For minimization, these quantities satisfy the standard inequality chain

$$z_{\text{ws}} \leq z_{\text{rp}} \leq z_{\text{eev}},$$

so EVPI and VSS can be interpreted as distances between implementable, perfect-information, and deterministic-approximation paradigms (Birge & Louveaux, 2011).

2.4.4 Nonconvexity and Integer Models

Many optimization problems in operations research involve discrete decisions (e.g., on/off, selection, assignment). These are commonly modeled with MILPs, which have a linear objective and linear constraints, but restrict some variables to be integer or binary. If the integrality restrictions are dropped, a linear program is obtained whose feasible set is a convex polyhedron. The nonconvexity in a MILP comes entirely from requiring integrality: the feasible set becomes a discrete collection of points (Nemhauser & Wolsey, 1988). This is a major source of computational hardness, but it also lets the model represent logical and combinatorial structure.

State-of-the-art MILP solvers rely on convex relaxations to compute bounds and then use search to enforce integrality. In a minimization problem, the LP relaxation gives a lower bound, while any feasible integer solution gives an upper bound. Branch-and-bound explores a tree of subproblems defined by partial assignments of integer variables, with the goal of closing the gap between the best bound and the best known feasible solution. Cutting planes further strengthen the relaxation by adding valid inequalities that move it closer to the convex hull of integer-feasible solutions (Nemhauser & Wolsey, 1988). Solvers typically stop when the remaining optimality gap falls below a user-specified tolerance, reflecting practical time limits and numerical precision. That is, MILP solvers report a deterministic optimality gap: the difference between the best relaxation-based bound and the best incumbent feasible solution (Nemhauser & Wolsey, 1988).

In stochastic programs with integer decisions, finite-scenario methods produce deterministic equivalents that are often large MILPs, with variables indexed by scenario and constraints that link scenarios through first-stage decisions. If recourse decisions are also integer-restricted, then the recourse value function is generally nonconvex and less

well-behaved, which weakens classical convex decomposition arguments and motivates specialized decomposition and bounding methods (Carøe & Tind, 1998). More generally, combining uncertainty with integrality increases computational difficulty, consistent with known hardness results in stochastic programming (Dyer & Stougie, 2006).

2.4.5 Non-Anticipativity for Multistage Decisions

A multistage model is only meaningful if decisions can be implemented sequentially as uncertainty unfolds. This is enforced via *non-anticipativity*: at stage t, decisions may depend on past and present observations, but not on future realizations (Dupacová et al., 2000).

This non-anticipativity requires that the stage-t decision x_t is measurable with respect to the information available just before (or at) stage t. Using the notation of Dupacová et al. (2000), if $\mathcal{F}_{t-1}$ is the σ-field generated by observations up to stage $t-1$, then

$$x_t \text{ is } \mathcal{F}_{t-1}\text{-measurable (equivalently, } \mathcal{F}_{t-1}\text{-adapted).} \tag{2.11}$$

Intuitively, two scenarios that are indistinguishable up to stage t must induce the same stage-t decision.

In scenario-tree formulations, non-anticipativity is enforced by equality constraints across scenarios that share the same history up to the relevant stage. An explicit example is provided by Ding et al. (2018), who enforce that the first-stage decisions are identical across all scenarios and, more generally, that stage-t decisions must coincide whenever scenarios share identical histories up to stage $t-1$. Using their notation with scenario indices i, j and history descriptors ϑ, representative constraints are:

$$x_{1,i}(\vartheta_{1,i}) = x_{1,j}(\vartheta_{1,j}), \qquad \forall i, j, \tag{2.12}$$

$$x_{t,i}(\vartheta_{t,i}) = x_{t,j}(\vartheta_{t,j}), \qquad \forall i, j \text{ with } (\vartheta_{1:i}, \ldots, \vartheta_{t-1,i}) = (\vartheta_{1,j}, \ldots, \vartheta_{t-1,j}). \tag{2.13}$$

These constraints operationalize the measurability condition in a finite scenario tree and ensure the resulting solution corresponds to an implementable policy (Ding et al., 2018; Dupacová et al., 2000).

2.4.6 Curse of Dimensionality

The practical challenge in multistage stochastic programming is that discretizing uncertainty by scenario trees can lead to extremely large deterministic equivalents. If uncertainty sources are independent, the number of scenarios required to represent their joint effect typically grows multiplicatively, i.e., as the product of the number of scenarios per source (Dupacová et al., 2000). This phenomenon is often referred to as the *curse of dimensionality* (or curse of problem size) and is a central reason why sampling, decomposition, and approximation methods are needed.

From a theoretical viewpoint, hardness results further justify approximation. Dyer and Stougie (2006) show that, under natural assumptions, broad classes of two-stage stochastic programs are computationally hard (including #P-hardness for evaluating the recourse expectation in certain settings) and that certain multistage variants can be PSPACE-hard.

These results provide formal support for the widespread observation that direct deterministic equivalent solution approaches become infeasible as the number of stages, scenarios, or decision couplings increases (Dyer & Stougie, 2006).

2.4.7 Sample Average Approximation

A standard approach to handle expectations in stochastic programming is SAA, which replaces the true expectation by an empirical mean over a finite MC sample (Shapiro, 2011; Shapiro, 2007).

Let $\xi^1, \ldots, \xi^N$ be i.i.d. samples from the distribution of ξ. The SAA objective for a two-stage problem is

$$\min_{x \in X} \hat{f}_N(x) := \frac{1}{N} \sum_{i=1}^{N} F(x, \xi^i), \tag{2.14}$$

which yields an approximate solution by solving a deterministic optimization problem built from sampled scenarios (Shapiro, 2007).

The theoretical appeal of SAA is that, under mild regularity conditions, $\hat{f}_N(x)$ converges to $f(x) = \mathbb{E}[F(x, \xi)]$ as $N \to \infty$ (pointwise by the law of large numbers), and the optimal value of the SAA problem converges to the true optimal value (Shapiro, 2007).

In multistage settings, a common workflow is to first generate a random sample (or sampled scenario tree) and then apply a decomposition method (e.g., SDDP for linear multistage problems) to the resulting SAA instance; Shapiro (2011) analyzes convergence properties precisely in this *SAA-then-solve* framework.

2.4.8 Optimality Gap Estimation

Let the true (risk-neutral) objective be $h(x) := \mathbb{E}[h(x, \xi)]$, and let $x^\star$ denote an optimal solution of the true problem. For any candidate solution $\hat{x}$, the (true) optimality gap is

$$G(\hat{x}) := h(\hat{x}) - h(x^\star) \geq 0.$$

In stochastic programming, expectations are rarely available in closed form, so $h(\hat{x})$ and $h(x^\star)$ cannot typically be evaluated exactly. Consequently, assessing solution quality requires statistical estimation procedures that provide (i) an estimate of $G(\hat{x})$ and (ii) uncertainty quantification via confidence intervals.

A widely used approach is the MC bounding procedure of Mak et al. (1999), which constructs a gap estimator using independent replications. For replication $r = 1, \ldots, n$, one forms a paired difference

$$G_r := \hat{z}_r^{\mathrm{UB}} - \hat{z}_r^{\mathrm{LB}},$$

where $\hat{z}_r^{\mathrm{UB}}$ is an unbiased MC estimate of the candidate policy value $h(\hat{x})$ (obtained by simulating $\hat{x}$ on an evaluation sample), and $\hat{z}_r^{\mathrm{LB}}$ is a statistical lower bound on the true optimum value $h(x^\star)$ derived from solving an SAA instance (or from replication-based bounding) (Mak et al., 1999). The paired structure reduces variance by correlating the upper- and lower-bound components within each replication.

The estimated gap is the sample mean

$$\hat{G} := \frac{1}{n} \sum_{r=1}^{n} G_r, \qquad \hat{\sigma}_G^2 := \frac{1}{n-1} \sum_{r=1}^{n} (G_r - \hat{G})^2.$$

Assuming G_r are i.i.d. with finite variance, an approximate $(1-\alpha)$-confidence interval for $G(\hat{x})$ is

$$\left[\hat{G} - t_{1-\alpha/2, n-1} \frac{\hat{\sigma}_G}{\sqrt{n}},\ \hat{G} + t_{1-\alpha/2, n-1} \frac{\hat{\sigma}_G}{\sqrt{n}} \right],$$

where $t_{1-\alpha/2, n-1}$ is the appropriate Student-t quantile (Mak et al., 1999).

2.4.9 Scenario Generation and Reduction

Many applications involve continuous or high-dimensional uncertainty models (e.g., forecast errors or copula-based demand processes). To solve stochastic programs numerically, these continuous distributions are commonly approximated by discrete scenario sets or scenario trees (Dupacová et al., 2000; Shapiro, 2011). In a multistage setting, the scenario tree represents possible future evolutions of the stochastic process; each root-to-leaf path corresponds to a full scenario trajectory (Dupacová et al., 2000).

Dupacová et al. (2000) emphasize that scenario trees should be representative relative to the optimization problem being modeled. They discuss multiple construction principles, including clustering of simulated paths, importance sampling, and moment-matching

approaches, and highlight that design choices (branching schemes, distance measures, weights) materially affect solution quality and computational tractability (Dupacová et al., 2000). This motivates validating the quality of a discrete approximation against the underlying stochastic model.

Because the number of scenarios needed to capture joint uncertainty can grow rapidly (Dupacová et al., 2000), scenario reduction seeks a smaller discrete measure (or smaller tree) that remains close to the original distribution in a probability-metric sense. A canonical formulation is: given an empirical/scenario measure P supported on a large set, find a measure supported on a subset of prescribed cardinality that minimizes a distance to P (Dupačová et al., 2003; Heitsch & Römisch, 2009). Forward and backward selection algorithms for such optimal (or approximately optimal) reduction problems are developed by Heitsch and Römisch (2009).

A widely used probability metric for comparing discrete approximations to a target distribution is the Wasserstein distance (Earth Mover's Distance for $p = 1$, Euclidean for $p = 2$). Given probability measures μ and ν on a metric space $(\mathcal{X}, d)$, the p-Wasserstein distance is

$$W_p(\mu, \nu) := \left(\inf_{\pi \in \Pi(\mu,\nu)} \int_{\mathcal{X}\times\mathcal{X}} d(x, y)^p \, \mathrm{d}\pi(x, y) \right)^{1/p}, \tag{2.15}$$

where $\Pi(\mu, \nu)$ is the set of couplings with marginals μ and ν (Peyré & Cuturi, 2020; Villani, 2009).

2.4.10 Computational Approaches

Stochastic MILPs are typically formed by combining a deterministic MILP with a finite set of scenarios representing uncertainty. The resulting problem size grows roughly linearly with the number of scenarios and exponentially with the number of stages, making naive solution approaches impractical (Apap & Grossmann, 2017; Torres Figueroa et al., 2022). As a result, two broad computational strategies are commonly used: solving the extensive form (EF) directly, or applying decomposition-based methods like progressive hedging (PH).

The EF is the deterministic equivalent MILP obtained by introducing a copy of the recourse variables for each scenario and enforcing non-anticipativity constraints (NACs) to ensure that decisions made before uncertainty is revealed are identical across scenarios (Watson et al., 2012). This yields a single, large MILP that can be solved using a general-purpose solver such as CPLEX or Gurobi.

For problems with a modest number of scenarios and stages, the EF is often the simplest and most reliable approach (Torres Figueroa et al., 2022). However, as the number of scenarios or stages increases, the EF quickly becomes computationally intractable due to the growth in variables, constraints, and branching complexity (Apap & Grossmann, 2017).

PH is a decomposition-based algorithm inspired by augmented Lagrangian and ADMM methods. PH relaxes the NACs and adds quadratic penalty terms that discourage disagreement among scenario-specific solutions (Boland et al., 2017; Watson et al., 2012). This relaxation allows the problem to decompose by scenario, so that one MILP per scenario is solved independently at each iteration. Scenario solutions are then coordinated by updating multipliers and penalties to drive the decisions toward non-anticipative consensus.

For convex stochastic programs, PH is provably convergent. In the mixed-integer setting, however, the problem is nonconvex, and standard PH becomes a heuristic: it often produces high-quality feasible solutions and sometimes useful bounds, but it does not guarantee convergence to a globally optimal solution (Boland et al., 2017; Watson et al., 2012). Despite this limitation, PH has been successfully applied to large-scale stochastic MILPs in areas such as power systems and logistics, particularly because it naturally supports parallel computation (Torres Figueroa et al., 2022).

In practice, EF formulations are typically preferred for small to medium-sized instances due to their simplicity and robustness (Torres Figueroa et al., 2022). For problems with

many scenarios or stages, decomposition-based methods become essential. PH is a powerful heuristic in this regime.

2.5 Forecasting and Reliability

To operationalize a stochastic optimization model, the underlying probability distributions of exogenous drivers must be characterized. This section covers the statistical methods for demand forecasting and reliability modeling used to generate the scenario trajectories that represent future operational uncertainty in the case study. Additionally, it connects reliability-based lifetime models to the generation of product return flows, necessary for closing the loop in this supply chain model.

2.5.1 Demand Forecasting

Demand forecasting provides a quantitative basis for supply chain planning and scenario development by translating historical patterns and contextual drivers of demand into expectations about future requirements. In the context of supply chain optimization, forecasts serve as exogenous inputs that shape decisions on capacity, inventory, sourcing, and risk mitigation; the work by Simchi-Levi et al. (2022) dedicates many sections to this context. Accurate demand forecasts may reduce planning uncertainty, limit undesirable effects such as the bullwhip effect (Lee et al., 1997), and enable the evaluation of alternative future states under both normal and disrupted conditions.

A common complication in practice is intermittent demand, characterized by many zero-demand periods interspersed with non-zero demands that may also be highly variable. Two widely discussed subtypes are lumpy demand (high intermittency and high variability of non-zero sizes) and erratic demand (lower intermittency but high variability) (A. A. Syntetos et al., 2005). Because intermittency and size variability represent distinct sources of uncertainty, it is useful to quantify them separately when selecting an appropriate forecasting model.

A standard diagnostic is the Syntetos–Boylan demand classification, based on (i) the average demand interval (ADI) and (ii) the squared coefficient of variation of non-zero demand sizes (CV^2) (A. A. Syntetos et al., 2005). Let $\{y_t\}_{t=1}^{n}$ denote demand observed over n periods, and let $m = |\{t : y_t > 0\}|$ be the number of non-zero periods. The ADI is

$$\mathrm{ADI} = \frac{n}{m},$$

and variability in positive demand sizes is summarized by

$$\mathrm{CV}^2 = \left(\frac{s^+}{\bar{y}^+}\right)^2,$$

where $\bar{y}^+$ and s^+ are the sample mean and standard deviation computed over $\{y_t : y_t > 0\}$. Using the thresholds of A. A. Syntetos et al. (2005), demand is commonly categorized as: (i) smooth if $\mathrm{ADI} < 1.32$ and $\mathrm{CV}^2 < 0.49$; (ii) intermittent if $\mathrm{ADI} \geq 1.32$ and $\mathrm{CV}^2 < 0.49$; (iii) erratic if $\mathrm{ADI} < 1.32$ and $\mathrm{CV}^2 \geq 0.49$; and (iv) lumpy if $\mathrm{ADI} \geq 1.32$ and $\mathrm{CV}^2 \geq 0.49$. This taxonomy is useful because lumpy demand reflects uncertainty both in occurrence and in magnitude conditional on occurrence; in such settings, single-component time-series models can perform poorly, motivating two-component approaches that model occurrence and size separately (e.g., Croston-type methods and related variants) (A. Syntetos et al., 2015).

Forecasting in asset- and service-intensive organizations should distinguish between routine operational demand and irregular, project-driven demand. In many contexts, such as higher education institutions (HEIs), operational activities generate recurrent requirements (e.g., campus services, facilities upkeep, and transport), whereas capital works and major maintenance arise sporadically due to asset aging, regulatory change, expansion, or strategic priorities (Leal Filho et al., 2015). This structure yields long

stretches of zero demand punctuated by concentrated events. Consequently, the processes governing demand occurrence and demand magnitude may differ, motivating intermittent-demand models over continuous time-series approaches.

In order to go about forecasting, historical demand records are required appropriate temporal granularity, and, for intermittent demand, the retention of zero-demand periods. Depending on the context, relevant exogenous variables such as prices, promotions, inventory signals, and seasonal or environmental factors may need assessment (Fildes et al., 2009; Gooijer & Hyndman, 2006). The preparation of such data takes typically place by removal of duplicates, treatment of missing values, consistency checks across systems, and evaluation of anomalies to distinguish genuine demand shifts from errors (Bandara et al., 2020). The selection of a historical date range, depends on the stability of supply-chain structures, data-collection systems, and external conditions; older data may represent circumstances that no longer apply (Seyedan & Mafakheri, 2020). A chosen window must therefore contain enough observations, whether seasonal cycles or intermittent-demand events, while remaining representative of current operations.

Consequently, modeling these components separately may improve accuracy and interpretability (Rožanec & Mladenić, 2021; Türkmen et al., 2021). For the occurrence process, inter-arrival times are typically modeled using geometric or negative binomial distributions in discrete time, or exponential distributions in continuous time, capturing the stochastic nature of whether demand occurs in a given period (A. Syntetos et al., 2015; Türkmen et al., 2021). When demand occurs, its size is commonly modeled using distributions such as the normal, gamma, lognormal, or Poisson, depending on whether order quantities are symmetric, right-skewed, continuous, or discrete, with zero-inflated or hurdle variants used where overdispersion or excess zeros must be accommodated (A. Syntetos et al., 2015; Türkmen et al., 2021).

2.5.2 Dependence Modeling

Many multivariate models must represent dependence: how random variables move together, beyond what can be explained by their individual marginal distributions. Copulas provide a standard way to do this. Sklar's theorem states that any joint distribution function F with marginals $F_1, \ldots, F_d$ can be written as

$$F(x_1, \ldots, x_d) \;=\; C\big(F_1(x_1), \ldots, F_d(x_d)\big), \tag{2.16}$$

where $C : [0,1]^d \to [0,1]$ is a copula that captures the dependence structure. For continuous marginals, the copula C is unique (Joe, 2014).

A common parametric choice is the Gaussian copula. It is defined by taking a multivariate normal distribution with correlation matrix Σ^{corr} and using its dependence structure together with arbitrary marginals. Concretely, for a random vector $X = (X_1, \ldots, X_d)$ with marginals $F_1, \ldots, F_d$, the Gaussian-copula construction can be written as

$$U_i = F_i(X_i), \qquad Y_i = \Phi^{-1}(U_i), \qquad Y \sim \mathcal{N}(0, \Sigma^{\text{corr}}), \tag{2.17}$$

so dependence is imposed in the latent Gaussian vector Y while the original marginals are retained via the inverse transforms. Equivalently, sampling can be described by drawing $Y \sim \mathcal{N}(0, \Sigma^{\text{corr}})$, setting $U_i = \Phi(Y_i)$, and mapping back using

$$X_i = F_i^{-1}(U_i) = F_i^{-1}\big(\Phi(Y_i)\big), \qquad i = 1, \ldots, d. \tag{2.18}$$

Gaussian copulas are popular because they are simple and computationally convenient, but they are limited in representing certain types of extremal dependence (e.g., strong tail dependence) (Joe, 2014).

In computational settings, dependence also affects geometry: when variables are correlated, Euclidean distances can be dominated by high-variance directions and correlation structure, which can distort distance-based procedures. A standard remedy is whitening (decorrelation). Given a random vector X with covariance matrix Σ, a whitening transform applies a matrix S such that

$$\tilde{X} = SX \quad \text{and} \quad \text{Cov}(\tilde{X}) = S\Sigma S^\top = I, \tag{2.19}$$

so the transformed variables have unit variance and are uncorrelated. For example, one common choice is $S = \Sigma^{-1/2}$ when Σ is positive definite (Kessy et al., 2015). Whitening goes beyond standardization (which rescales variables but leaves correlations intact) and often yields more stable behavior for distance-based clustering and approximation in correlated spaces (Kessy et al., 2015).

2.5.3 Reliability Modeling

Reliability modeling describes the random time until a unit fails, is replaced, or (more generally) exits use. Let $L \geq 0$ denote a lifetime (time-to-failure). The key objects are the cumulative distribution function (CDF) $F(t) = \mathbb{P}(L \leq t)$ and the survival function $S(t) = \mathbb{P}(L > t) = 1 - F(t)$.

A widely used lifetime model is the Weibull distribution due to its flexibility and simple closed forms (McCool, 2012; Weibull, 1951). With shape parameter $\beta > 0$ and scale parameter $\eta > 0$, the Weibull CDF, probability distribution function (PDF), and survival function are

$$F(t; \beta, \eta) = 1 - \exp\left(-\left(\frac{t}{\eta}\right)^{\beta}\right), \qquad t \geq 0, \tag{2.20}$$

$$f(t; \beta, \eta) = \frac{\beta}{\eta}\left(\frac{t}{\eta}\right)^{\beta-1} \exp\left(-\left(\frac{t}{\eta}\right)^{\beta}\right), \tag{2.21}$$

$$S(t; \beta, \eta) = \exp\left(-\left(\frac{t}{\eta}\right)^{\beta}\right). \tag{2.22}$$

The hazard (failure rate) summarizes the instantaneous failure risk conditional on survival up to time t:

$$h(t) \;=\; \frac{f(t)}{S(t)} \;=\; \frac{\beta}{\eta}\left(\frac{t}{\eta}\right)^{\beta-1}. \tag{2.23}$$

This expression makes the role of β explicit: $\beta < 1$ implies a decreasing hazard (early-life failures), $\beta = 1$ yields a constant hazard (the exponential special case), and $\beta > 1$ implies an increasing hazard (wear-out/aging) (Jiang & Murthy, 2011).

Many planning and control problems operate on a discrete time grid (e.g., monthly decisions). A continuous lifetime model can be converted into interval event probabilities by integrating over a time bin. For an interval length $\Delta > 0$, the probability that failure occurs in $(t, t + \Delta]$ is

$$\mathbb{P}(t < L \leq t + \Delta) \;=\; F(t + \Delta) - F(t). \tag{2.24}$$

Often the relevant quantity is the conditional probability of failing during the next interval given survival up to t:

$$\mathbb{P}(t < L \leq t + \Delta \mid L > t) \;=\; \frac{F(t + \Delta) - F(t)}{S(t)} \;=\; 1 - \frac{S(t + \Delta)}{S(t)}. \tag{2.25}$$

For Weibull lifetimes, using (2.22) gives the closed form

$$\mathbb{P}(t < L \leq t + \Delta \mid L > t) \;=\; 1 - \exp\left(-\frac{(t + \Delta)^{\beta} - t^{\beta}}{\eta^{\beta}}\right). \tag{2.26}$$

This provides a direct bridge between continuous-time reliability assumptions and discrete-time transition probabilities used in multi-period models.

Finally, event occurrences over time are often described using renewal processes. A (simple) renewal process is defined by i.i.d. inter-arrival times $X_1, X_2, \ldots$ (e.g., successive lifetimes), with event times $T_n = \sum_{i=1}^{n} X_i$ and counting process $N(t) = \max\{n : T_n \leq t\}$ (Jiang, 2020). Renewal models are a standard way to connect lifetime distributions to sequences of replacement/return events when units are repeatedly renewed after each event.

Chapter 3

Methodology

This chapter details the methodology of the multistage stochastic life cycle optimization framework. It integrates matrix-based LCA endogenously within a supply chain optimization model to support decisions under joint inventory and operational uncertainty. The methodology follows a hierarchical structure: (i) defining the network topology and process-flow relationships (Sections 3.1–3.3); (ii) characterizing uncertainty in life cycle impacts, demand, and returns (Sections 3.4–3.6); and (iii) specifying the computational solution and statistical evaluation procedures (Sections 3.7–3.8).

3.1 Supply Chain Topology

The system under study is modeled as a network of interconnected processes and flows spanning a product life cycle, from upstream resource acquisition to end-of-life treatment. Figure 3.1 illustrates the generic closed-loop supply chain topology of an arbitrary supplier considered in this work. The topology is partitioned into seven sections (I–VII), each corresponding to a distinct life cycle phase. Their operational scopes are summarized in Table 3.1.

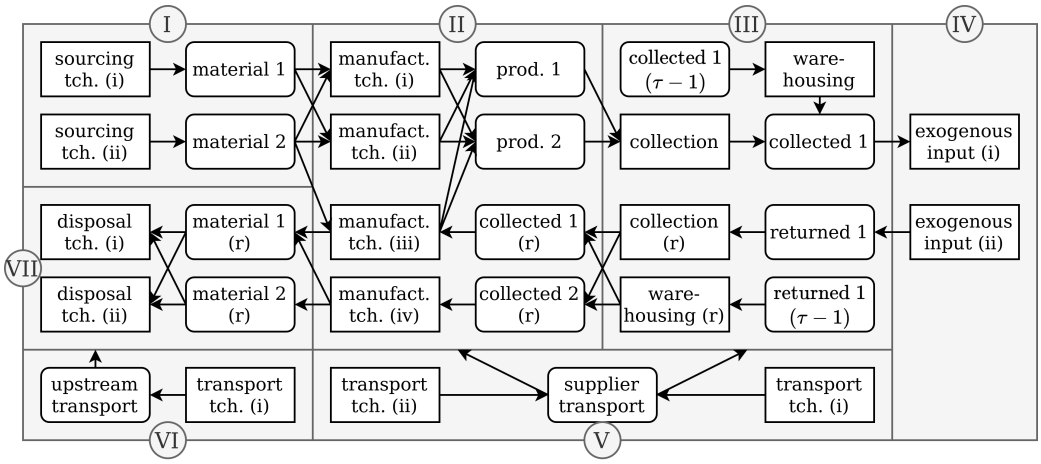

Figure 3.1: Schematic representation of the generic supply chain topology of an arbitrary supplier, illustrating material flows between upstream acquisition, manufacturing, use, and end-of-life stages. Processes, represented as technologies (tch.), are illustrated with square corners and flows with rounded corners. Numerals I–VII correspond to the phases defined in Table 3.1.

Processes (i.e. technologies) are represented as nodes with square corners, while flows of products and services are represented by nodes with rounded corners. Both forward and reverse material movements are captured within the same network structure, allowing

recovered materials from reverse flows to re-enter forward production stages and thereby enabling closed-loop interactions.

The topology supports flexible routing and transformation of flows through multiple alternative processes. Different processes can satisfy the same functional demand, representing technological substitution in manufacturing, transportation, or end-of-life treatment. Processes may also generate multiple output flows, capturing co-production phenomena such as disassembly and material separation. Shared infrastructure, including warehousing and transportation, is modeled through common capacity usage by multiple flows and processes, introducing structural coupling across supply chain stages.

Reverse flows are embedded symmetrically, allowing returned products to be directed either to disposal or to higher-value recovery options. Intermediate collection and storage nodes decouple stochastic return timing from downstream processing capacities. All flows are time-indexed, such that each flow is defined for every discrete time stage considered in the model, ensuring consistent representation of material balances, capacity constraints, and technology selection decisions over the supply chain lifecycle.

Table 3.1: Phases and sub-processes of the supply chain topology.

#	Phase	Description
I	Raw materials	Extraction, sourcing, and pre-processing of materials used as inputs to supplier manufacturing. Includes upstream (higher-tier) supplier activities prior to the focal supplier.
II	Manufacturing	Supplier production of goods and services using relevant manufacturing technologies. Incorporates virgin and circular inputs (e.g., recovered materials) and may generate residuals destined for disposal or treatment.
III	Distribution	Collection, handling, and warehousing of products, including inspection, consolidation, long-term storage, and management of reverse flows.
IV	Use phase	Downstream use that determines exogenous demand and drives upstream production and distribution. May also generate exogenous reverse flows (e.g., returns or post-consumer collection).
V	Transportation (supplier)	Supplier-managed transport supporting manufacturing and distribution, potentially involving multiple transport modes and technologies.
VI	Transportation (upstream)	Transport associated with upstream sourcing activities and downstream disposal or treatment of materials.
VII	End-of-life	Collection and treatment of waste streams, including recycling, recovery, and final disposal of products and materials.

3.2 Mathematical Formalization

The formalization of the supply chain topology (Figure 3.1) builds upon and extends the flow-based model proposed by Hülagü et al. (2025). First, the process–flow relationships are algebraically defined using the technology matrix $\mathbf{A}$ and intervention matrix $\mathbf{B}$, following the computational structure defined in Section 2.1.1. Each column of $\mathbf{A}$ corresponds to a process (technology), while each row corresponds to a flow. Flows entering a process are represented by negative entries, and flows leaving a process by positive entries. When a flow can be produced by multiple technologies, the corresponding row contains multiple positive entries across different columns. Conversely, a single process may generate

multiple output flows, resulting in several positive entries within the same column.

As discussed in Section 2.1.1, this construction generally yields a non-square matrix $\mathbf{A}$, which is therefore not invertible. In the present optimization framework, consistent with the linear programming formulation, no inversion of $\mathbf{A}$ is required: process activity levels are decision variables, and material balances are enforced directly through linear constraints. As a result, no explicit allocation rule for distributing supply across competing technologies is needed.

Each process is associated with environmental and economic exchanges, which are represented by the intervention matrix $\mathbf{B}_\tau \in \mathbb{R}^{N_e \times N_p}$. The impacts of these exchanges are characterized through the characterization matrix $\mathbf{Q} \in \mathbb{R}^{N_i \times N_e}$, where $N_i \in \mathbb{N}$ is the number of impact categories. Exogenous inputs to the system, such as final demand, are represented by the flow vector $\mathbf{f}_\tau \in \mathbb{R}^{N_f}$. The scaling vector $\mathbf{s}_\tau \in \mathbb{R}^{N_p}$ determines the activity level of each process at stage $\tau \in \{1, \dots, T\}$ with $T \in \mathbb{N}$ the number of time stages. To allow for more production than required, i.e. abundance of flows, an additional vector $\mathbf{e} \in \{0, 1\}^{N_f}$ is defined specifying whether overflow is possible (1) or not (0). All matrices and vectors are indexed by time where appropriate. In order to move flows between time stages, a Leslie matrix $\mathbf{K} \in \mathbb{R}^{N_f \times N_f}$ is defined that carries flows from $\tau - 1$ to τ.

3.2.1 Decision variables

The decision variables consist of the supplier choice $x \in \mathcal{X}$, where $\mathcal{X}$ is a finite set of candidate suppliers, and the sequence of recourse integer decisions $\mathbf{s}_\tau(\cdot) \in \mathbb{N}_0^{N_p}$ over the planning horizon. A finite horizon with stages $\tau = 1, \dots, T$ is considered. At the initial stage $\tau = 1$, a single supplier is selected from the candidate set $\mathcal{X}$. This here-and-now decision determines both the distribution of the supplier-specific stochastic intervention matrix $\tilde{\mathbf{B}}(x) \in \mathbb{R}^{N_e \times N_p}$ and the deterministic capacity vector $\mathbf{u}(x) \in \mathbb{R}_{\geq 0}^{N_c}$ (with $N_c \in \mathbb{N}$ the number of capacity constraints) available in subsequent stages.

3.2.2 Stochastic variables

At the beginning of stage $\tau = 2$, the random matrix $\tilde{\mathbf{B}}(x)$ is realized and remains fixed for all stages $\tau = 2, \dots, T$. At each stage $\tau = 2, \dots, T$, the random demand vector $\tilde{\mathbf{f}}_\tau \in \mathbb{R}^{N_f}$ is revealed. The information revealed at stage τ is collected in the random vector $\xi_\tau(x)$, and the history of realized uncertainty up to stage τ is denoted by $\xi_{2:\tau}(x) \coloneqq (\xi_2(x), \dots, \xi_\tau(x))$. The random vector $\xi_\tau(x)$ is defined as

$$\xi_\tau(x) \coloneqq \begin{cases} (\tilde{\mathbf{B}}(x), \tilde{\mathbf{f}}_2), & \tau = 2, \\ \tilde{\mathbf{f}}_\tau, & \tau = 3, \dots, T. \end{cases} \tag{3.1}$$

At each stage $\tau = 2, \dots, T$, the system responds to the realized information by selecting a recourse decision $\mathbf{s}_\tau\big(\xi_{2:\tau}(x)\big)$. This induces a random inventory vector $\tilde{\mathbf{g}}_\tau \in \mathbb{R}^{N_e}$ given by

$$\tilde{\mathbf{g}}_\tau\big(x, \xi_{2:\tau}(x)\big) = \tilde{\mathbf{B}}(x)\, \mathbf{s}_\tau\big(\xi_{2:\tau}(x)\big), \qquad \tau = 2, \dots, T. \tag{3.2}$$

The inventory vector is partitioned into environmental and economic components. The same block structure applies to the intervention matrix:

$$\tilde{\mathbf{g}}_\tau\big(x, \xi_{2:\tau}(x)\big) = \begin{bmatrix} \tilde{\mathbf{g}}_\tau^{\text{env}}\big(x, \xi_{2:\tau}(x)\big) \\ \tilde{\mathbf{g}}_\tau^{\text{ecn}}\big(x, \xi_{2:\tau}(x)\big) \end{bmatrix}, \qquad \tilde{\mathbf{B}}(x) = \begin{bmatrix} \tilde{\mathbf{B}}^{\text{env}}(x) \\ \tilde{\mathbf{B}}^{\text{ecn}}(x) \end{bmatrix}.$$

3.2.3 Life cycle impact variables

Environmental and economic impacts are calculated using the characterization operator $\mathbf{Q}$ and aggregated by the weighting operator $\mathbf{W}(\lambda) \in \mathbb{R}^{1 \times N_i}$, where $\lambda \in [0, 1]$ controls the trade-off between the two dimensions. A weighted objective is chosen here over other multi-objective techniques (such as ϵ-constraint) because it allows for a tractable and

feasible exploration of the efficient frontier in a stochastic setting. Their block structure is given by

$$\mathbf{Q} = \begin{bmatrix} \mathbf{Q}^{\text{env}} & \mathbf{0} \\ \mathbf{0} & \mathbf{Q}^{\text{ecn}} \end{bmatrix}, \qquad \mathbf{W}(\lambda) = \begin{bmatrix} (1-\lambda)\,\mathbf{W}^{\text{env}} & \lambda\,\mathbf{W}^{\text{ecn}} \end{bmatrix}, \tag{3.3}$$

where $\mathbf{W}^{\text{env}} \in \mathbb{R}^{1\times N_i^{\text{env}}}$ and $\mathbf{W}^{\text{ecn}} \in \mathbb{R}^{1\times N_i^{\text{ecn}}}$ are row vectors weighting the environmental and economic impact categories respectively, with $N_i^{\text{env}} \in \mathbb{N}$ and $N_i^{\text{ecn}} \in \mathbb{N}$ denoting the number of environmental and economic impact categories such that $N_i = N_i^{\text{env}} + N_i^{\text{ecn}}$.

The expected aggregate impact for supplier x and trade-off parameter λ is defined as the accumulated weighted impact over the planning horizon:

$$h(x;\lambda) = \mathbb{E}\left[\sum_{\tau=2}^{T} \mathbf{W}(\lambda)\,\mathbf{Q}\,\tilde{\mathbf{g}}_\tau\big(x, \xi_{2:\tau}(x)\big)\right]. \tag{3.4}$$

Due to the multistage structure and continuous uncertainty, the expectation in (3.4) cannot be evaluated in closed form. It is therefore approximated numerically using SAA, as described in Section 2.4.7.

3.2.4 Objective Function and Constraints

For a fixed value of the weighting parameter λ, the optimal supplier choice $x^\star(\lambda)$ and the associated recourse decisions are obtained as the solution of a multistage stochastic mixed-integer program. The objective function aggregates environmental and economic impacts through the weighted cost function $h(x;\lambda)$, while feasibility is ensured by structural constraints.

Formally, the problem is given by

$$\min_{x\in\mathcal{X},\,\mathbf{s}_\tau(\cdot)} \quad h(x;\lambda) \tag{3.5}$$

$$\text{s.t.} \quad \operatorname{diag}(\mathbf{e})\Big(\mathbf{A}_\tau \mathbf{s}_\tau(\xi_{2:\tau}(x)) - \tilde{\mathbf{f}}_\tau\Big) + \mathbf{K}\mathbf{A}_{\tau-1}\mathbf{s}_{\tau-1}(\xi_{2:\tau-1}(x)) = \mathbf{0}, \tag{3.6}$$

$$\operatorname{diag}(\mathbf{1}-\mathbf{e})\Big(\mathbf{A}_\tau \mathbf{s}_\tau(\xi_{2:\tau}(x)) - \tilde{\mathbf{f}}_\tau\Big) + \mathbf{K}\mathbf{A}_{\tau-1}\mathbf{s}_{\tau-1}(\xi_{2:\tau-1}(x)) \geq \mathbf{0}, \tag{3.7}$$

$$\mathbf{A}_\tau^{\text{out}}\,\mathbf{s}_\tau(\xi_{2:\tau}(x)) \leq \mathbf{u}(x), \tag{3.8}$$

for all stages $\tau = 2,\dots,T$.

The feasible set is defined by four interrelated classes of constraints: non-anticipativity, flow balance, inventory carry-over, and process capacity. Together, they ensure physical consistency over time, respect the information structure of the stochastic process, and adheres to supplier-specific technological limits.

Non-anticipativity

Recourse decisions must respect the temporal structure of information revelation; see Section 2.4.5. At each stage $\tau = 2,\dots,T$, the process-activity vector $\mathbf{s}_\tau$ may depend only on uncertainty realized up to that stage, summarized by the history $\xi_{2:\tau}(x)$. In particular, decisions cannot anticipate future realizations of stochastic parameters.

Flow balance and inventory carry-over

Material conservation is enforced at every stage through flow-balance constraints derived from the technology matrix $\mathbf{A}_\tau$. These constraints link production decisions, realized demand $\tilde{\mathbf{f}}_\tau$, and inventory transferred from the previous stage.

The vector $\mathbf{e}$ specifies which flows must satisfy demand exactly and which may be produced in excess. For flows with $e_i = 1$, equality constraint (3.6) enforces exact balance.

For flows with $e_i = 0$, inequality constraint (3.7) prevents shortages while allowing surplus generation.

Intertemporal coupling is modeled through inventory carry-over via the matrix $\mathbf{K}$. The term $\mathbf{K}\mathbf{A}_{\tau-1}\mathbf{s}_{\tau-1}$ represents the portion of previous-stage production that remains available at stage τ, ensuring intertemporal consistency without introducing explicit inventory state variables.

Capacity feasibility

Operational feasibility is enforced through capacity constraints on process activities. Let $\mathbf{A}_\tau^{\text{out}}$ denote the submatrix of $\mathbf{A}_\tau$ corresponding to output flows that consume limited infrastructure or processing resources. For each stage τ, activity levels must satisfy constraint (3.8), where $\mathbf{u}(x)$ is the supplier-specific capacity vector determined by the first-stage decision x.

Because $\mathbf{u}(x)$ is fixed once the supplier is selected, these constraints couple the strategic supplier decision with all future recourse actions, linking long-term selection with operational feasibility.

3.3 Process Capacity Modeling

The optimization model enforces demand satisfaction at every stage and does not include backlogging. As a consequence, feasibility depends on whether upstream technologies can deliver the required flows within the planning period. In classical supply-chain formulations, lead time is often handled via ordering policies, safety stocks, and inventory/backlog dynamics. To address the strict no-backlog requirement in the flow-based formulation used here, a heuristic capacity estimation is proposed where lead time is interpreted as an indicator of congestion and translated into an effective deliverable capacity per planning period. This yields a supplier- and technology-specific upper bound on deliverable flows, captured by the capacity vector $\mathbf{u}(x)$ in the process-capacity constraint (3.8).

3.3.1 From lead time to effective capacity

Let p index a deliverable category (or constrained flow group) and let t index historical planning periods (e.g., weeks or months). Define the realized throughput in period t as the total delivered quantity

$$\mathrm{TP}_{p,t} \;=\; \sum_{i \in \mathcal{I}_{p,t}} q_i,$$

where $\mathcal{I}_{p,t}$ is the set of delivered orders in category p during period t, and q_i is the delivered quantity of order i.

Lead time is measured as elapsed time between order placement and delivery. To reduce the influence of small orders, a quantity-weighted lead time is computed per period:

$$\mathrm{LT}_{p,t} \;=\; \frac{\sum_{i \in \mathcal{I}_{p,t}} q_i \,\mathrm{lt}_i}{\sum_{i \in \mathcal{I}_{p,t}} q_i}, \tag{3.9}$$

where lt_i is the realized lead time of order i.

Under capacity pressure, queueing effects imply that lead time increases with throughput and typically exhibits a nonlinear *knee* as throughput approaches a maximum sustainable level. This motivates estimating an effective capacity $\hat{C}_p$ from historical $(\mathrm{TP}_{p,t}, \mathrm{LT}_{p,t})$ observations. Two estimation routes are considered.

Knee-based capacity

If the throughput–lead-time relationship shows a clear transition into a congested regime, $\hat{C}_p$ is defined as the throughput level at the knee of the empirical curve (optionally after smoothing to reduce measurement noise). The resulting $\hat{C}_p$ is interpreted as the maximum throughput that can be sustained without rapidly increasing lead times.

Quantile-based capacity

If a robust knee cannot be identified (e.g., weak congestion signals, limited observations, or noisy lead-time measurements), a conservative nonparametric proxy is used:

$$\hat{C}_p \;=\; Q_\alpha(\mathrm{TP}_{p,t})\,, \qquad \alpha \in (0,1), \tag{3.10}$$

where $Q_\alpha(\cdot)$ denotes the α-quantile across historical periods. This defines capacity as a high but typical historically achieved throughput level and is robust to incidental high lead-times.

3.3.2 Feasibility considerations

Because the formulation enforces demand satisfaction without backlogging, overly restrictive capacity bounds can render the optimization problem infeasible. In practice, this is addressed by ensuring that (i) $\mathbf{u}(x)$ is calibrated to the model's time scale and units, and (ii) the set of feasible suppliers/technologies includes at least one option with sufficient capacity for all relevant demand scenarios. If infeasibility remains a concern, an additional modeling layer is required (e.g., explicit slack for unmet demand with a penalty, alternative recourse actions, or a redesign of the scenario set), since infeasibility may not be resolved through anticipation.

3.4 LCIA Uncertainty Quantification

This section describes how process-level environmental and economic impact coefficients are constructed, how their uncertainty is quantified, and how it is propagated to obtain the stochastic process matrix used in the optimization model. The overall procedure is schematically illustrated in Figure 3.2.

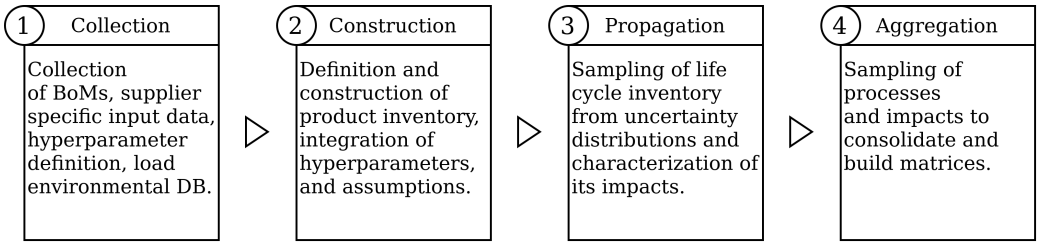

Figure 3.2: Schematic step-by-step overview of the impact modeling and uncertainty propagation procedure: background data collection (1), construction and (hyper)parameterization of the inventory (2), propagation of uncertainty by means of MC samples (3), and aggregation to consolidate different processes to build matrices for scenario generation (4).

Background information is collected from environmental databases, product disclosures, industry documentation, and expert input. Process models may depend on uncertain global parameters (e.g., material densities, recycling rates), differ by composition or variant (e.g., color variants). Collected information are then used to construct inventory processes. Subsequently, the inventory processes are characterized to individual functional unit product processes. Subsequently, these products are sampled to include probabilistic assumptions regarding exchanges or process occurrence. The result is a $\tilde{\mathbf{P}}$ consisting of $\mathbf{A}$ and $\tilde{\mathbf{B}}(x)$ that is structurally characterized with $\mathbf{Q}$ when propagated.

3.4.1 Process Matrix Construction

The characterized process matrix $\tilde{\mathbf{P}}$ is composed of a technology matrix and stochastic characterized inventories. The technology matrix is restricted to unit entries (-1, 0, or 1)

and therefore encodes only the structural relationships between processes. All quantitative scaling, such as material quantities or transport demand, is embedded in the intervention inventories (e.g., a requirement of 10 tonne-kilometers is represented as a single transport activity, with the corresponding impact encoded in its characterization). As a result, matrix size is limited and the stochasticity of LCI is captured entirely in $\tilde{\mathbf{B}}(x)$.

Foreground unit processes follow the supply chain topology (Figure 3.2). Raw material processes are aggregated to reduce the size of the technology matrix. The consolidated process matrix $\mathbf{A}$ is then used as input to the optimization algorithm.

3.4.2 Quantification and Sampling Procedure

LCI and LCIA calculation is implemented in Brightway, a Python framework for life cycle assessment (Mutel, 2017), enabling efficient recomputation of LCI and LCIA results under perturbed inputs. Prior to uncertainty propagation, parameter expressions are parsed and registered in Brightway's parameter system. Uncertainty distributions for global parameters are represented using `stats_arrays` (Mutel, 2025), while uncertainty in background LCI exchanges follows the ecoinvent-specified pedigree indices.

Uncertainty is modeled in selected foreground exchanges (e.g., chance of necessary repair), translated as background processes, and propagated through inventory calculation and impact assessment via MC simulation. Characterization factors for LCIAs are treated nominally. While LCIA factors also involve methodological and scenario uncertainty, they are not addressed here to keep the uncertainty model focused on inventory imprecision and to maintain computational tractability (Heijungs & Huijbregts, 2004; Huijbregts, 1998). Consistent with the limitations of the pedigree matrix discussed in Section 2.2.1, the uncertainties here reflect epistemic data quality rather than pure empirical process variability. Consequently, the resulting impact distributions should be interpreted as conditional on the data-quality model.

Let N denote the number of MC iterations. In each iteration, a coherent set of uncertain process-data realizations and exchange coefficients are drawn, the LCI is recomputed for each foreground unit process, and the corresponding LCIA scores are evaluated. The sample size N governs the stability of estimated moments and tail quantities (e.g., quantiles) of the induced impact distributions. Here, $N = 10{,}000$ iterations are used, consistent with common practice in LCA uncertainty propagation when pedigree-based uncertainties dominate (Heijungs, 2019; Igos et al., 2018).

Economic impacts are also treated stochastically with probabilistic choices introduced before variants are aggregated (e.g., multiple compositions or color variants mapped to a single product flow that the optimization model cannot distinguish). In such cases, the induced variation is treated analogously and propagated together with environmental impacts within the same MC framework.

3.5 Demand Scenario Generation

This section presents a distribution-based method for generating stochastic demand scenarios suitable for multistage inventory and procurement planning under intermittent and lumpy demand. The objective is to produce scenario sets that (i) reproduce intermittency (zero-demand periods), (ii) capture tailed order quantities, (iii) preserve seasonal structure when present, and (iv) allow dependence across product categories. The method is intentionally generic; the subsequent case study instantiates the framework by selecting a particular temporal granularity, a two-component intermittent-demand specification, and a dependence-induction procedure.

3.5.1 Rationale and modeling requirements

Intermittent demand is characterized by many zero-demand periods and sporadic, often large orders. Classical intermittent-demand forecasting methods such as Croston and its variants (e.g., SBA) primarily provide mean forecasts and do not yield full predictive

distributions, limiting their usefulness for scenario generation, tail-risk evaluation, and service-level assessment (Teunter & Duncan, 2009). In contrast, parametric distributional modeling specifies complete distributions for order timing and order size, supports extrapolation beyond observed values, and remains feasible under sparse data (Rožanec & Mladenić, 2021; A. Syntetos et al., 2015). Bootstrapping is a possible alternative but cannot generate unseen extremes and degrades under limited history (Teunter & Duncan, 2009). Simulation-based approaches for intermittent demand address these limitations by generating full demand distributions and are therefore well suited as stochastic inputs for optimization (Willemain et al., 2004).

3.5.2 Two-component intermittent-demand model

To represent intermittency explicitly, demand is modeled using a two-component decomposition (occurrence and size). This approach addresses the specific challenges of 'lumpy' demand identified in Section 2.5.1, where classical single-point forecasts fail to capture tail risks (Babaï et al., 2014; A. Syntetos & Boylan, 2005). Let demand for product category k be observed on a base time grid (e.g., weeks). Define (i) an occurrence process that determines when orders occur, and (ii) a size process that determines the order quantity conditional on occurrence.

Occurrence component

Order occurrences are modeled as a renewal process in terms of inter-arrival times (gaps) between non-zero demand periods. For each category k, inter-arrival times are fitted to candidate parametric distributions (exponential, geometric, and Weibull). Parameters are estimated using standard likelihood-based methods, and candidate fits are compared using graphical diagnostics and Akaike information criterion (AIC).

Size component

Conditional on an occurrence, strictly positive order sizes are modeled independently using established parametric families for lumpy demand: lognormal, gamma, and zero-truncated negative binomial distributions (Rožanec & Mladenić, 2021; A. Syntetos et al., 2015). As with the occurrence model, parameters are estimated by likelihood methods and models are selected based on AIC and visual fit. The baseline assumption is that occurrence and size are approximately independent; this assumption must be evaluated empirically in a case-specific context using dependence diagnostics between gaps and order sizes.

3.5.3 Temporal granularity and seasonal structure

Scenario generation must align with the decision stages of the optimization model while retaining enough observations for reliable estimation. The framework therefore distinguishes between: (i) a fitting grid used for model estimation (chosen to maximize statistical power), and (ii) a decision grid used for planning and scenario reporting.

In the present study, model estimation is performed on a weekly grid to exploit the larger sample size typical in intermittent-demand settings, while scenarios are aggregated and used at monthly resolution to match planning practice, economic impacts, and improve the stability of decision-stage inputs (A. Syntetos & Boylan, 2005; Türkmen et al., 2021). Aggregation is performed during simulation by generating order occurrences and sizes on the fitting grid and then summing to monthly totals.

Seasonality is incorporated when diagnostics indicate systematic month-of-year differences. Because intermittent-demand seasonality commonly manifests more strongly in demand volumes than in arrival frequencies (Babaï et al., 2014; Rožanec & Mladenić, 2021; A. Syntetos et al., 2015), the default seasonal mechanism used here is month-of-year stratification of the size distribution: each calendar month $m \in \{1, \ldots, 12\}$ is assigned its own size-distribution parameters for each product category. The occurrence process is retained as stationary unless arrival-seasonality is clearly indicated. This approach

preserves seasonal variation without requiring full time-series state-space modeling and remains feasible under sparse demand.

3.5.4 Cross-category dependence

When multiple product categories are planned jointly, ignoring cross-category dependence can distort portfolio risk and lead to under- or over-conservative procurement decisions. Dependence is therefore characterized empirically by aggregating historical demand to the decision grid (monthly totals) and computing a Pearson correlation matrix across categories.

To induce dependence across simulated category-wise marginals while retaining the fitted univariate distributions, a copula-based dependence mapping is used. Specifically, a Gaussian copula is constructed with correlation matrix $\mathbf{\Sigma}^{\text{corr}}$ estimated from historical monthly totals. Let $\mathbf{L}$ denote its Cholesky factor. Independent standard normal samples $\mathbf{Z}_\tau \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ are first correlated via $\mathbf{Y}_\tau = \mathbf{L}\mathbf{Z}_\tau$, converted to uniforms by $U_{j,\tau} = \Phi\big((\mathbf{Y}_\tau)_j\big)$, and then mapped to demand values using the calibrated inverse CDFs from (3.15):

$$\tilde{f}_{j,\tau} = \hat{F}^{-1}_{j,\tau}\big(\Phi\big((\mathbf{L}\mathbf{Z}_\tau)_j\big)\big)\,. \tag{3.11}$$

This procedure ensures that the simulated demands for different products move together in a way that reflects the correlations observed in the historical data, while maintaining the empirical marginal distributions.

3.5.5 Sampling Procedure

Demand scenarios are generated over a planning horizon of T decision periods (e.g., months). For each scenario and each product category, monthly demand is constructed in three steps. First, the timing of orders within the month is simulated by drawing inter-arrival times on the fitted grid, which determines the weeks in which orders occur. Second, an order size is sampled for each occurrence from the selected size distribution (including month-of-year stratification where applicable). Third, the sampled orders are aggregated to obtain the monthly demand total for that category. This procedure produces an initial set of monthly demand scenarios for each category, generated independently across categories. Cross-category dependence is then introduced by coupling these marginal scenarios using the procedure described in Section 3.5.4, resulting in joint demand scenarios.

The number of simulated scenarios is selected to yield stable estimates of tail quantities, such as high quantiles, that drive service-level and robustness considerations. In this work, 100,000 draws are generated at the marginal-simulation stage to stabilize upper-quantile behavior. The multistage optimization model is subsequently constructed using a reduced scenario representation derived from these samples (see Section 2.4.9). All simulations are executed with fixed random seeds and consistent sampling settings across categories.

3.6 Return Scenario Generation

Closing the supply chain loop requires an explicit representation of return flows. This section presents a method for generating stochastic product return scenarios for closed-loop planning when return timing is governed by product use and degradation rather than by an independent demand-like process. The approach models returns as end-of-use events from an installed base, using cohort-based reliability modeling. The subsequent case study instantiates this framework for furniture categories using Weibull lifetime distributions with assumed parameters due to limited return-specific reliability data.

3.6.1 Rationale of lifetime modeling

In addition to forward flows from manufacturing to the use phase, closed-loop planning requires a representation of use-to-return flows. In some settings, returns can be modeled

as occurring after a fixed post-use delay, which can support endogenous return modeling in certain closed-loop systems (Hülagü et al., 2025); this is not adopted here. For durable products, return timing is typically governed by usage intensity, degradation, and replacement decisions, and is therefore stochastic and heterogeneous across units. Accordingly, returns are represented through exogenous time-to-return distributions drawn used in reliability theory, analogous to demand as an external stochastic input.

3.6.2 Cohort-based reliability model

Returns are generated using a cohort-based framework in which historical sales (or deployments) form an installed base that ages over time. Historical purchase records are aggregated into cohorts indexed by product category k and delivery month τ, using delivery date as the start-of-use time. Let $Q_{k,\tau}$ denote the number of units delivered in cohort (k, τ). Each unit in cohort (k, τ) is assumed to have an independent time to return L_k drawn from a category-specific lifetime distribution.

The cohort perspective is essential for durable goods because future returns depend not only on a lifetime model but also on the age structure of the installed base. This structure is naturally represented by the set of active cohorts and their sizes.

3.6.3 Lifetime distribution and discretization

Time-to-return is modeled using a Weibull distribution (see Section 2.5.3) (Rinne, 2008; Weibull, 1951). This choice allows for flexible representation of hazard behaviors, specifically the wear-out phase ($\beta > 1$) relevant to durable furniture, as defined in Eq. 2.23 (Frolov et al., 2024; Gómez et al., 2023; McCool, 2012). For each category k, lifetime is modeled as

$$L_k \sim \text{Weibull}(\beta_k, \eta_k),$$

with shape β_k and scale η_k. In the absence of return/failure observations sufficient for statistical estimation, parameters may be set from engineering judgment, literature values, or durability expectations.

For scenario generation on a discrete planning grid (months), the continuous lifetime distribution is converted into age-specific monthly return probabilities. Let a denote the age in months. Define the discrete probability that a unit of age a returns within the next month as

$$p_k(a) = F_k(a + 1) - F_k(a),$$

where $F_k(\cdot)$ is the Weibull cumulative distribution function. For a forecast horizon of T months starting at month t_0, each cohort (k, τ) induces a probability vector over the horizon months, with residual probability mass representing returns occurring outside the forecast window.

This construction assumes that the likelihood of return increases with age when $\beta_k > 1$. While early-return phenomena (e.g., infant mortality, policy-driven take-backs, aesthetic replacement) may be relevant in other contexts and can be captured by alternative lifetime models or mixture distributions, the baseline assumption adopted here is that return propensity increases with time in use. This is considered a pragmatic representation for institutional purchasing where replacement cycles and approvals tend to reduce very early returns.

3.6.4 Sampling procedure

Given cohort sizes $Q_{k,\tau} \in \mathbb{N}_0$ and the corresponding discrete probability vector over horizon months, stochastic monthly returns are generated by multinomial sampling at the cohort level. Specifically, for each cohort (k, τ) a vector of counts is drawn that allocates the cohort's units across potential return months within the horizon, with an additional outside-horizon category absorbing remaining probability mass. Repeating this procedure yields independent realizations of return volumes that reflect stochastic lifetime outcomes.

Let $S \in \mathbb{N}$ denote the number of MC realizations. For each realization $s \in \{1, \ldots, S\}$, simulated returns are accumulated across cohorts to obtain category-level monthly returns $R_{k,t}^{(s)} \in \mathbb{N}_0$:

$$R_{k,t}^{(s)} = \sum_{\tau \leq t} R_{k,\tau \to t}^{(s)}, \qquad t = t_0, \ldots, t_0 + T - 1,$$

where $R_{k,\tau \to t}^{(s)} \in \mathbb{N}_0$ denotes returns in month t originating from cohort (k, τ) in realization s. The resulting scenario set $\{R_{k,t}^{(s)}\}$ provides a distribution over monthly return volumes for each category that can be used directly in downstream planning and optimization.

Simulations are executed using fixed random seeds and consistent sampling settings across categories to support reproducibility. The baseline model treats unit lifetimes as independent conditional on category-level parameters and treats categories as independent.

In principle, returns may be statistically dependent across categories (e.g., coordinated refurbishments, facility moves) or linked to demand through replacement behavior. When data permit, such dependencies can be introduced using the same dependence-induction mechanisms as in the demand scenario generator (e.g., copula-based coupling on monthly totals). In the present study, dependence modeling is limited by data availability and is therefore handled conservatively in the case study discussion.

3.7 Computational Approach

The multistage stochastic mixed-integer program introduced in the previous sections cannot be solved in closed form due to the presence of continuous uncertainty, expectation operators, and NAC requirements across stages. A computational strategy is therefore required to transform the model into a tractable equivalent that can be handled by optimization solvers, while still preserving the essential stochastic structure of the problem. This section describes the numerical methodology adopted for this purpose.

3.7.1 Sample Average Approximation

To solve the multistage stochastic mixed-integer program, the SAA method is employed (see Section 2.4.7). This approach approximates the continuous expectation in (3.4) using a finite set of scenarios organized into a scenario tree structure (Heitsch & Römisch, 2009). The tree explicitly enforces non-anticipativity by bundling scenarios with identical histories, thereby discretizing the information filtration. While SAA provides asymptotic convergence to the true optimal value (Kleywegt et al., 2002), solution quality depends on the representativeness of the generated scenario tree (Park & Hanasusanto, 2024).

To obtain a tractable deterministic formulation of the objective in (3.4), the continuous distribution of the uncertainty process is approximated by a discrete distribution supported on the scenario paths Ω of a scenario tree. Each path $\omega \in \Omega$ is assigned a probability $p_\omega > 0$ with $\sum_{\omega \in \Omega} p_\omega = 1$, induced by the underlying distributions and the tree-generation procedure. For each scenario $\omega \in \Omega$ and supplier $x \in \mathcal{X}$, the history of realized uncertainty up to stage τ is defined as in (3.12):

$$\xi_{2:\tau}(x, \omega) \coloneqq \big(\xi_2(x, \omega), \ldots, \xi_\tau(x, \omega)\big), \qquad \omega \in \Omega,\ \tau = 2, \ldots, T. \tag{3.12}$$

Using the corresponding recourse decisions $\mathbf{s}_\tau\big(\xi_{2:\tau}(x, \omega)\big)$ and inventories $\tilde{\mathbf{g}}_\tau\big(x, \xi_{2:\tau}(x, \omega)\big)$, the scenario-wise aggregate impact is given by (3.13):

$$h(x; \lambda, \omega) = \sum_{\tau=2}^{T} \mathbf{W}(\lambda)\, \mathbf{Q}\, \tilde{\mathbf{g}}_\tau\big(x, \xi_{2:\tau}(x, \omega)\big), \qquad \omega \in \Omega. \tag{3.13}$$

For a fixed scenario tree with path set Ω and associated probabilities $\{p_\omega\}_{\omega \in \Omega}$, the SAA approximation of (3.4) is then

$$\hat{h}(x; \lambda) = \sum_{\omega \in \Omega} p_\omega\, h(x; \lambda, \omega), \tag{3.14}$$

which defines a deterministic objective conditional on the chosen scenario tree and its probability weights. The special case $p_\omega = 1/|\Omega|$ for all $\omega \in \Omega$ corresponds to equally likely scenarios.

In the numerical experiments, the resulting deterministic mixed-integer programs with explicit scenario dependence, as defined by (3.14) and the constraints (3.6)-(3.8) instantiated for all $\omega \in \Omega$, are implemented and solved in Python using Pyomo (Bynum et al., 2021; Hart et al., 2011), MPI-SPPY (Knueven et al., 2023), and Gurobi (Gurobi Optimization, LLC, 2025).

3.7.2 Scenario Tree Construction

The scenario trees used in the SAA approximation are constructed from empirical distributions associated with the two sources of uncertainty in the model: the random intervention matrix $\tilde{\mathbf{B}}(x)$ and the stage-wise demand vectors $\tilde{\mathbf{f}}_\tau$. The construction is implemented using a Python port of the scenario-tree library Scentrees (Kirui et al., 2020; van Mastrigt, 2025); more information can be found in Appendix C.1.

Impact scenarios are constructed directly from MC simulation results in impact space. For each supplier x, samples of the environmental and economic impacts for all (unit) processes are obtained as $\tilde{\mathbf{H}}(x) = \mathbf{Q}\tilde{\mathbf{B}}(x)$. To obtain a representative set of scenarios, the impact columns are partitioned into empirical quantile bins of approximately equal probability mass, from each of which one observation is sampled. This produces the representative impact matrices $\widehat{\mathbf{H}}^{(1)}(x), \ldots, \widehat{\mathbf{H}}^{(M)}(x)$ for each supplier, jointly capturing environmental and economic impacts.

Dependence between products at a given stage is enforced using the Gaussian copula model defined in Section 3.5.4. The stage-wise demand scenarios are constructed from historical data for installed and returned products. For each component j of the demand vector and each stage τ, an empirical marginal distribution is estimated by smoothing the historical observations $\{f_{j,\tau}^{(m)}\}_m$ with a kernel density estimator (KDE) (Weglarczyk, 2018), computing the associated CDF $\hat{F}_{j,\tau}$, and obtaining its numerical inverse $\hat{F}_{j,\tau}^{-1}$. Samples from this marginal model are then generated as

$$\tilde{f}_{j,\tau} = \hat{F}_{j,\tau}^{-1}(U_{j,\tau}), \qquad U_{j,\tau} \sim \mathcal{U}(0,1), \tag{3.15}$$

which provides a nonparametric approximation of the empirical demand distribution at stage τ.

A number of independent demand trajectories is then generated from the copula model, each trajectory providing one simulated demand vector for every stage $\tau = 2, \ldots, T$, sampled independently across stages. Before these simulated trajectories are passed to a distance-based scenario-tree approximation routine, their components are standardized by means of a whitening transformation. As described in Section 2.5.2 (Equation 2.19), this removes linear correlations to produce an isotropic representation, ensuring that distance-based tree construction is not biased by high-variance directions (Kessy et al., 2015). Let $\widehat{\boldsymbol{\Sigma}}_f$ denote the empirical covariance matrix of all simulated demand vectors (pooled across stages and replications), and let $\mathbf{S}$ be a symmetric matrix satisfying $\mathbf{S}\,\widehat{\boldsymbol{\Sigma}}_f\,\mathbf{S}^\top \approx \mathbf{I}$. The whitened demand vectors are defined as

$$\tilde{\mathbf{f}}'_\tau = \mathbf{S}\,\tilde{\mathbf{f}}_\tau. \tag{3.16}$$

Without this preprocessing, the Euclidean distances used by the Scentrees algorithm are dominated by both high-volume products and strong correlations between products, causing the scenario tree to overrepresent these directions. Standardization removes only the scale effect, whereas whitening removes both scale and covariance distortions, reducing the risk that a single product or correlated group of products overwhelms the distance metric and yielding a more balanced multivariate approximation.

Given a user-specified set of branching factors, Scentrees constructs a multistage tree in the whitened space by adjusting node states and conditional probabilities so that the

distribution induced by the tree matches the empirical distribution of the whitened reference trajectories $\{\tilde{\mathbf{f}}'_\tau\}$. After the approximation procedure, node states are mapped back to the original correlated space using $\mathbf{S}^{-1}$ and rounded to integers to reflect the count nature of the demands and returns.

Finally, the joint scenario tree is obtained by combining the LCA scenarios with the leaf paths of the demand tree. Independence between $\tilde{\mathbf{H}}(x)$ and $(\tilde{\mathbf{f}}_\tau)_{\tau=2}^T$ implies that the probability of a joint scenario is the product of the probability of the LCA scenario and the probability associated with the corresponding demand path. For each such joint scenario, the uncertainty vectors $\xi_\tau(x)$ are recovered according to (3.1), with $\tilde{\mathbf{H}}(x)$ fixed along the path and $\tilde{\mathbf{f}}_\tau$ determined by the demand-tree nodes at stage τ. These scenario paths constitute the finite support Ω used in the SAA objective (3.14), and their tree structure enforces non-anticipativity of the recourse decisions $\mathbf{s}_\tau(\cdot)$ across all scenarios with identical histories.

3.8 Evaluation Framework

The computational approach yields implementable SAA-based approximations of the underlying multistage stochastic program, but the resulting objective values and decisions remain subject to both sampling variability (through process impact data, scenario generation and tree approximation) and optimization error (through non-convex mixed-integer solves). An explicit evaluation framework is therefore required to (i) quantify how reliable a candidate supplier decision is under uncertainty, (ii) assess whether modeling uncertainty and information structure materially improves performance, and (iii) disentangle which sources of uncertainty dominate the variability in the resulting impact. This section introduces three complementary analyses to address these needs: a statistically interpretable optimality-gap and confidence analysis based on independent scenario-tree replications, standard value-of-information metrics (EVPI and VSS) to benchmark the benefit of information and stochastic modeling, and a variance decomposition that apportions impact uncertainty to structural (tree), scenario-choice, and process-data components.

3.8.1 Gap and Confidence Analysis

To solve and evaluate the stochastic program, several independent scenario trees, and their associated objectives of the form (3.14), are generated in order to assess sampling variability and evaluate the quality of the resulting solutions. Each tree is constructed from independently drawn realizations of the random intervention matrix $\tilde{\mathbf{B}}(x)$, characterized with $\mathbf{Q}$, and the stage-wise demand vectors $(\tilde{\mathbf{f}}_\tau)_{\tau=2}^T$. On these trees, two complementary evaluations are performed. First, a fixed supplier $\hat{x} \in \mathcal{X}$ is evaluated by re-optimizing only the recourse decisions, which yields an estimator of the expected impact of the corresponding policy. Second, fully optimal (up to solver tolerances) SAA problems are solved on independent trees, which yields an estimator of the optimal objective value of the underlying stochastic program. The combination of these two evaluations allows the construction of a statistical estimate of the optimality gap of $\hat{x}$ in the sense of Mak-Morton-Wood (Mak et al., 1999).

Let $\hat{x} \in \mathcal{X}$ denote the supplier selected by solving one particular SAA problem or by some other design choice. For each independent scenario tree $r = 1, \ldots, M_{\mathrm{UB}}$ with path set $\Omega^{(r)}$ and probabilities $\{p_\omega^{(r)}\}_{\omega \in \Omega^{(r)}}$, the policy induced by $\hat{x}$ is evaluated by solving

$$\hat{z}_r^{\mathrm{UB}} = \min_{\mathbf{s}_\tau(\cdot)} \sum_{\omega \in \Omega^{(r)}} p_\omega^{(r)} \sum_{\tau=2}^{T} \mathbf{W}(\lambda)\, \mathbf{Q}\, \tilde{\mathbf{g}}_\tau\big(\hat{x}, \xi_{2:\tau}(\hat{x}, \omega)\big), \tag{3.17}$$

$$\bar{G}_n \pm t_{1-\alpha/2,\, n-1}\, \frac{s_G}{\sqrt{n}}, \tag{3.18}$$

subject to the balance and capacity constraints (3.6)-(3.8) instantiated for all $\tau = 2, \ldots, T$ and all $\omega \in \Omega^{(r)}$, with the supplier fixed to $\hat{x}$. In this problem, only the recourse decisions $\mathbf{s}_\tau(\cdot)$ are decision variables; the intervention matrix $\tilde{\mathbf{B}}(\hat{x})$ and capacity vector $\mathbf{u}(\hat{x})$

are determined by the chosen supplier. The resulting values $\{\hat{z}_r^{\mathrm{UB}}\}_{r=1}^{M_{\mathrm{UB}}}$ are independent realizations of a MC estimator of the expected aggregate impact $h(\hat{x};\lambda)$ defined in (3.4).

To obtain a benchmark for the optimal performance of the stochastic program, a second set of independent scenario trees $r = 1, \ldots, M_{\mathrm{LB}}$ is generated. On each such tree, the full SAA problem is solved without fixing the supplier choice. For a scenario tree with path set $\Omega^{(r)}$ and probabilities $\{p_\omega^{(r)}\}_{\omega\in\Omega^{(r)}}$, this problem takes the form

$$\hat{z}_r^{\mathrm{LB}} = \min_{x\in\mathcal{X},\,\mathbf{s}_\tau(\cdot)} \sum_{\omega\in\Omega^{(r)}} p_\omega^{(r)}\, h\big(x;\lambda,\omega\big), \tag{3.19}$$

again subject to the constraints (3.6)-(3.8) for all stages and scenarios. Because the underlying deterministic equivalents are mixed-integer and therefore non-convex, the values $\hat{z}_r^{\mathrm{LB}}$ correspond to optimal or near-optimal solutions obtained by a MILP solver under the chosen tolerances. They are interpreted as approximate realizations of the optimal SAA value on each tree.

The optimality gap associated with the candidate supplier $\hat{x}$ is defined as

$$G = h(\hat{x};\lambda) - h(x^\star(\lambda);\lambda) \;\geq 0. \tag{3.20}$$

The optimality gap G is estimated using the paired-differences estimator of Mak, Morton and Wood (Mak et al., 1999), as outlined in Section 2.4.8. In the case where the same collection of scenario trees is used for both evaluations, that is, when $M_{\mathrm{UB}} = M_{\mathrm{LB}} = n$ and each pair $(\hat{z}_r^{\mathrm{UB}}, \hat{z}_r^{\mathrm{LB}})$ is computed on the same tree, the paired difference is $G_r = \hat{z}_r^{\mathrm{UB}} - \hat{z}_r^{\mathrm{LB}}$. The estimated gap $\hat{G}$ is given by the sample mean $\bar{G}_n$. The associated variance and standard error are computed as per standard paired-sample statistics, allowing for the construction of a $(1-\alpha)$-confidence interval using the Student-t distribution with $n-1$ degrees of freedom.

In addition to the absolute gap $\hat{G}$, a relative gap is reported whenever $\hat{\mu}_{\mathrm{LB}} > 0$,

$$\hat{G}_{\mathrm{rel}} = 100\,\frac{\hat{G}}{\hat{\mu}_{\mathrm{LB}}}\,\%, \tag{3.21}$$

which quantifies the estimated suboptimality of the candidate supplier $\hat{x}$ relative to the estimated optimal objective level. Taken together, the estimates $\hat{\mu}_{\mathrm{UB}}$, $\hat{\mu}_{\mathrm{LB}}$, the gap $\hat{G}$, and its confidence interval provide a statistically interpretable assessment of how close the decision $\hat{x}$ is to optimality, while explicitly acknowledging both sampling error from the SAA approximation and optimization error due to the non-convex mixed-integer nature of the underlying deterministic problems.

3.8.2 Value of Information

Beyond evaluating a single candidate decision by means of the optimality-gap analysis in Section 3.8.1, of interest is to quantify how much the decision maker could gain from improved information and from explicitly modeling uncertainty. This is done by computing the EVPI and the VSS. Both quantities are defined in terms of the objective values of suitably modified versions of the optimization problem in (3.5)-(3.8). Throughout this section, the trade-off parameter λ is treated as fixed, and all problem variants are mixed-integer and therefore non-convex. In practice, their objective values are obtained from optimal or near-optimal solutions returned by a MILP solver under the same tolerances as in the main analysis.

The reference problem is the multistage stochastic program already introduced in Section 3.7. For a fixed λ, its optimal value is

$$z_{\mathrm{rp}}(\lambda) = \min_{x\in\mathcal{X},\,\mathbf{s}_\tau(\cdot)} h(x;\lambda), \tag{3.22}$$

where $h(x;\lambda)$ is the expected aggregate impact defined in (3.4), and the minimization is taken over all suppliers $x \in \mathcal{X}$ and all non-anticipative recourse policies $\mathbf{s}_\tau(\cdot)$ satisfying (3.6)-(3.8). The value $z_{\mathrm{rp}}(\lambda)$ is denoted by RP and represents the best achievable

expected performance when decisions must be taken sequentially under uncertainty and obey non-anticipativity.

A first comparison is made with an idealized "wait-and-see" or perfect-information setting. In this case, the decision maker is assumed to know the entire realization of the uncertainty before any decision is taken, so that both the supplier choice and all recourse decisions may depend on the full scenario path $\omega \in \Omega$. For a given scenario ω, the corresponding perfect-information problem is

$$z_{\mathrm{ws}}(\lambda, \omega) = \min_{x(\omega) \in \mathcal{X},\, \mathbf{s}_\tau(\omega)} h\big(x(\omega); \lambda, \omega\big), \tag{3.23}$$

where the scenario-wise impact $h(x; \lambda, \omega)$ is given by (3.13), and non-anticipativity constraints across scenarios are dropped. The expected value of perfect information is then

$$z_{\mathrm{ws}}(\lambda) = \mathbb{E}\big[z_{\mathrm{ws}}(\lambda, \omega)\big] = \sum_{\omega \in \Omega} p_\omega \, z_{\mathrm{ws}}(\lambda, \omega), \tag{3.24}$$

with p_ω denoting the scenario probabilities on the underlying scenario tree. Comparing (3.22) and (3.24), the EVPI is defined consistent with the standard stochastic programming metrics reviewed in Section 2.4.3:

$$\mathrm{EVPI}(\lambda) = z_{\mathrm{rp}}(\lambda) - z_{\mathrm{ws}}(\lambda) \; \geq 0. \tag{3.25}$$

This quantity measures the expected reduction in impact that would be achievable if the entire uncertainty were revealed before any decision is made. In the LCA context considered here, $\mathrm{EVPI}(\lambda)$ quantifies the opportunity cost, in units of the weighted environmental-economic impact, of making supplier and recourse decisions under partial information.

A second comparison is made between the fully stochastic formulation and a simplified model in which all random quantities are replaced by their expectations. To this end, the expected intervention matrix and demand vectors are defined componentwise as

$$\bar{\mathbf{B}}(x) = \mathbb{E}\big[\tilde{\mathbf{B}}(x)\big], \qquad \bar{\mathbf{f}}_\tau = \mathbb{E}\big[\tilde{\mathbf{f}}_\tau\big], \quad \tau = 2, \ldots, T, \tag{3.26}$$

and collected in an expected parameter vector $\bar{\xi}(x)$. Substituting (3.26) into the balance and capacity constraints and into the impact computation yields a deterministic expected-value problem

$$z_{\mathrm{ev}}(\lambda) = \min_{x \in \mathcal{X},\, \mathbf{s}_\tau} h\big(x; \lambda, \bar{\xi}(x)\big), \tag{3.27}$$

whose formulation is structurally identical to (3.5)-(3.8) but with all randomness replaced by its expectation and with a single scenario path. Let $x_{\mathrm{ev}}(\lambda)$ denote a minimizer of (3.27). The corresponding recourse decisions need not be retained, since the expected-value problem is used only to define the first-stage decision $x_{\mathrm{ev}}(\lambda)$.

To evaluate how this deterministic decision performs under the full stochastic model, the expected-value solution is re-embedded into the original multistage program by fixing the supplier choice to $x_{\mathrm{ev}}(\lambda)$ and re-optimizing the recourse decisions under uncertainty. This leads to the expected result of the expected-value solution,

$$z_{\mathrm{eev}}(\lambda) = \min_{\mathbf{s}_\tau(\cdot)} h\big(x_{\mathrm{ev}}(\lambda); \lambda\big), \tag{3.28}$$

where the minimization is again taken over all non-anticipative recourse policies on the original scenario tree. The Value of the Stochastic Solution is defined as

$$\mathrm{VSS}(\lambda) = z_{\mathrm{eev}}(\lambda) - z_{\mathrm{rp}}(\lambda) \; \geq 0. \tag{3.29}$$

A positive value of $\mathrm{VSS}(\lambda)$ indicates that explicitly modeling uncertainty and enforcing non-anticipativity, as in (3.22), yields a better expected performance than optimizing against expected parameters alone and then using recourse decisions to adjust.

In practice, the quantities $z_{\mathrm{rp}}(\lambda)$, $z_{\mathrm{ws}}(\lambda)$, $z_{\mathrm{ev}}(\lambda)$, and $z_{\mathrm{eev}}(\lambda)$ cannot be computed exactly. Instead, they are approximated by SAA on independently generated scenario trees,

in the same fashion as for the optimality-gap analysis in Section 3.8.1. Let $r = 1, \dots, R$ index independent scenario trees, each with path set $\Omega^{(r)}$ and probabilities $\{p_\omega^{(r)}\}_{\omega \in \Omega^{(r)}}$. On each replication r, the SAA counterparts of (3.22), (3.24), (3.27), and (3.28) are solved, yielding approximate values

$$\hat{z}_{\mathrm{rp},r}(\lambda), \qquad \hat{z}_{\mathrm{ws},r}(\lambda), \qquad \hat{z}_{\mathrm{ev},r}(\lambda), \qquad \hat{z}_{\mathrm{eev},r}(\lambda),$$

from which replication-wise estimates of EVPI and VSS are formed as

$$\widehat{\mathrm{EVPI}}_r(\lambda) = \hat{z}_{\mathrm{rp},r}(\lambda) - \hat{z}_{\mathrm{ws},r}(\lambda), \qquad \widehat{\mathrm{VSS}}_r(\lambda) = \hat{z}_{\mathrm{eev},r}(\lambda) - \hat{z}_{\mathrm{rp},r}(\lambda). \tag{3.30}$$

The sample means

$$\overline{\mathrm{EVPI}}(\lambda) = \frac{1}{R}\sum_{r=1}^{R} \widehat{\mathrm{EVPI}}_r(\lambda), \qquad \overline{\mathrm{VSS}}(\lambda) = \frac{1}{R}\sum_{r=1}^{R} \widehat{\mathrm{VSS}}_r(\lambda), \tag{3.31}$$

provide MC estimators of $\mathrm{EVPI}(\lambda)$ and $\mathrm{VSS}(\lambda)$. Their sampling variances are estimated by

$$\hat{\sigma}^2_{\mathrm{EVPI}}(\lambda) = \frac{1}{R-1}\sum_{r=1}^{R}\big(\widehat{\mathrm{EVPI}}_r(\lambda) - \overline{\mathrm{EVPI}}(\lambda)\big)^2, \tag{3.32}$$

$$\hat{\sigma}^2_{\mathrm{VSS}}(\lambda) = \frac{1}{R-1}\sum_{r=1}^{R}\big(\widehat{\mathrm{VSS}}_r(\lambda) - \overline{\mathrm{VSS}}(\lambda)\big)^2, \tag{3.33}$$

leading to standard errors

$$\widehat{\mathrm{SE}}\big(\overline{\mathrm{EVPI}}(\lambda)\big) = \frac{\hat{\sigma}_{\mathrm{EVPI}}(\lambda)}{\sqrt{R}}, \qquad \widehat{\mathrm{SE}}\big(\overline{\mathrm{VSS}}(\lambda)\big) = \frac{\hat{\sigma}_{\mathrm{VSS}}(\lambda)}{\sqrt{R}}. \tag{3.34}$$

Under standard conditions on the SAA approximation and for sufficiently large R, the central limit theorem implies that these sample means are approximately normally distributed. As a consequence, confidence intervals for $\mathrm{EVPI}(\lambda)$ and $\mathrm{VSS}(\lambda)$ can be constructed analogously to (3.18) by using the estimated standard errors in (3.34) and a Student-t quantile with $R-1$ degrees of freedom.

3.8.3 Variance Decomposition

In this section, the aggregate impact of a fixed supplier $x \in \mathcal{X}$ and trade-off parameter $\lambda \in [0,1]$ is treated as a random quantity, and its uncertainty is decomposed into three components: (i) structural uncertainty arising from the construction of the scenario trees, (ii) scenario-choice uncertainty within a given tree (containing a set of decisions), and (iii) input variable uncertainty in the process-level LCI data. The analysis operationalizes the UA framework discussed in Section 2.2.3. By applying the law of total variance (Equation 2.6) to the linear LCA structure, the aggregate impact variance is decomposed into these three components.

Linear Combination of Process Impacts

Fix a supplier $x \in \mathcal{X}$ and a trade-off parameter $\lambda \in [0,1]$. For each scenario tree, indexed by $r \in \mathcal{R} := \{1, \dots, R\}$, let $\Omega^{(r)}$ denote the finite set of scenario paths (leaf paths) with associated probabilities $\{p_\omega^{(r)}\}_{\omega \in \Omega^{(r)}}$, as in the SAA formulation in Section 3.7.1. For a given tree r and path $\omega \in \Omega^{(r)}$, the scenario-wise aggregate impact is

$$h(x; \lambda, \omega, r) = \sum_{\tau=2}^{T} \mathbf{W}(\lambda)\, \mathbf{Q}\, \tilde{\mathbf{g}}_\tau\big(x, \xi_{2:\tau}(x, \omega, r)\big), \tag{3.35}$$

where the inventory vectors $\tilde{\mathbf{g}}_\tau\big(x, \xi_{2:\tau}(x, \omega, r)\big)$ are defined as in (3.2) along the path ω of tree r.

Because the underlying LCA model is linear, the scalar quantity in (3.35) can be written, for each pair (r, ω),as a linear combination of process-level impact factors. For a fixed supplier x, let P denote the number of distinct unit processes in the combined foreground-background system. For each stage $\tau = 2, \ldots, T$ and each scenario path $\omega \in \Omega^{(r)}$, the recourse decision $\mathbf{s}_\tau\big(\xi_{2:\tau}(x, \omega, r)\big)$ is a vector of process scaler levels, as introduced in (3.6)-(3.8).

Define the stage-aggregated activity vector for the entire scenario path (r, ω) as

$$\mathbf{s}_{r,\omega}(x) \;\coloneqq\; \sum_{\tau=2}^{T} \mathbf{s}_\tau\big(\xi_{2:\tau}(x, \omega, r)\big) \;\in\; \mathbb{R}^P, \tag{3.36}$$

whose i-th component represents the total activity of process i over all stages along the path (r, ω).

Furthermore, let

$$\tilde{\boldsymbol{\eta}}(x) = \big(\tilde{\eta}_1(x), \ldots, \tilde{\eta}_P(x)\big)^\top \tag{3.37}$$

denote a random vector of process-level impact factors, where $\tilde{\eta}_i(x)$ is the environmental or economic impact per unit activity of process i after transformation by $\mathbf{Q}$ and $\mathbf{W}(\lambda)$. The randomness in $\tilde{\boldsymbol{\eta}}(x)$ represents parameter uncertainty in the LCI database and is approximated by MC samples.

With these definitions, the scenario-wise aggregate impact admits the linear representation

$$\tilde{h}(x; \lambda, \omega, r) = \mathbf{s}_{r,\omega}(x)^\top \tilde{\boldsymbol{\eta}}(x), \tag{3.38}$$

which is the basis for the hierarchical uncertainty analysis and the variance decomposition that follows.

Representation of Uncertainty

Three sources of uncertainty are considered. The first one is structural (tree) uncertainty originating from the scenario-tree construction procedure. This yields R alternative trees, indexed by $r \in \mathcal{R}$, which approximate the underlying stochastic process in different ways. The index of the tree used for evaluation is modeled as a discrete random variable $\tilde{R}$ taking values in $\mathcal{R}$ with equal probabilities,

$$\mathbb{P}(\tilde{R} = r) = \pi_r \coloneqq \frac{1}{R}, \qquad r \in \mathcal{R}. \tag{3.39}$$

Second is the scenario-choice (activity) uncertainty originating from within trees. This is conditional on a given tree r, the scenario path ω is modeled as a discrete random variable $\tilde{\omega}$ taking values in $\Omega^{(r)}$ with probabilities

$$\mathbb{P}(\tilde{\omega} = \omega \mid \tilde{R} = r) = p_\omega^{(r)}, \qquad \omega \in \Omega^{(r)},\ r \in \mathcal{R}, \tag{3.40}$$

where $p_\omega^{(r)} > 0$ and $\sum_{\omega \in \Omega^{(r)}} p_\omega^{(r)} = 1$.

Lastly is the process-data uncertainty affecting the LCI impact factors, modeled as a random vector $\tilde{\boldsymbol{\eta}}(x)$ with an unknown distribution. In practice, this distribution is represented by MC samples. The analysis assumes that process-level impact factors are mutually independent across processes, so that the covariance matrix of $\tilde{\boldsymbol{\eta}}(x)$ is diagonal. This assumption isolates the contribution of individual unit-process variances but neglects cross-correlations inherent in the background database.

$$\big\{\tilde{\boldsymbol{\eta}}^{(k)}(x)\big\}_{k=1}^{K}, \qquad \tilde{\boldsymbol{\eta}}^{(k)}(x) = \big(\tilde{\eta}_1^{(k)}(x), \ldots, \tilde{\eta}_P^{(k)}(x)\big)^\top, \tag{3.41}$$

where each index k corresponds to a draw over all processes.

Combining these elements, the aggregate impact of supplier x at trade-off parameter λ is represented as the scalar random variable

$$\tilde{H}(x; \lambda) \coloneqq \tilde{h}\big(x; \lambda, \tilde{\omega}, \tilde{R}\big) = \mathbf{s}_{\tilde{R},\tilde{\omega}}(x, \lambda)^\top \tilde{\boldsymbol{\eta}}(x), \tag{3.42}$$

whose distribution reflects the combined effect of structural, scenario-choice, and process-data uncertainty. The corresponding expectation $\mathbb{E}[\tilde{H}(x; \lambda)]$ and variance $\mathrm{Var}[\tilde{H}(x; \lambda)]$ are approximated by MC simulation.

Monte Carlo Approximation

To approximate the distribution of $\tilde{H}(x;\lambda)$, an all-at-a-time MC simulation is performed. Let N denote the number of MC iterations. For each iteration $\ell = 1, \ldots, N$, a tree index is first sampled according to

$$R^{(\ell)} \sim \text{Categorical}(\pi_1, \ldots, \pi_R),$$

and, conditional on $R^{(\ell)} = r$, a scenario path is sampled as

$$\omega^{(\ell)} \sim \text{Categorical}\big(p_\omega^{(r)}\big)_{\omega \in \Omega^{(r)}}.$$

A realization of process-level impacts is then obtained. Consistent with the independence assumption, process impacts are sampled independently for each process i:

$$K_i^{(\ell)} \sim \{1, \ldots, K\} \text{ uniformly}, \quad \forall i = 1, \ldots, P$$

$$\tilde{\boldsymbol{\eta}}^{(\ell)}(x) = \big(\tilde{\eta}_1^{(K_1^{(\ell)})}(x), \ldots, \tilde{\eta}_P^{(K_P^{(\ell)})}(x)\big)^\top.$$

The deterministic scaling vector corresponding to the selected tree-scenario pair is subsequently retrieved,

$$\mathbf{s}^{(\ell)}(x,\lambda) \coloneqq \mathbf{s}_{R^{(\ell)},\omega^{(\ell)}}(x,\lambda).$$

Finally, the aggregate impact for replication ℓ is computed as

$$\tilde{H}^{(\ell)}(x;\lambda) = \big(\mathbf{s}^{(\ell)}(x,\lambda)\big)^\top \tilde{\boldsymbol{\eta}}^{(\ell)}(x).$$

The sample $\{\tilde{H}^{(\ell)}(x;\lambda)\}_{\ell=1}^N$ provides an empirical approximation of the distribution of (3.42). The corresponding sample mean and variance are

$$\hat{\mu}_H(x;\lambda) = \frac{1}{N}\sum_{\ell=1}^{N} \tilde{H}^{(\ell)}(x;\lambda), \tag{3.43}$$

$$\widehat{\text{Var}}\big[\tilde{H}(x;\lambda)\big] = \frac{1}{N-1}\sum_{\ell=1}^{N}\big(\tilde{H}^{(\ell)}(x;\lambda) - \hat{\mu}_H(x;\lambda)\big)^2. \tag{3.44}$$

From these, summary measures such as the coefficient of variation can be obtained as $\widehat{\text{CV}} = \sqrt{\widehat{\text{Var}}[\tilde{H}(x;\lambda)]}/\hat{\mu}_H(x;\lambda)$.

Three-Level Variance Decomposition

Beyond quantifying the overall uncertainty in $\tilde{H}(x;\lambda)$, it is of interest to apportion the variance to the three sources of uncertainty. To this end, the law of total variance is applied twice, with respect to the random indices $\tilde{R}$ and $\tilde{\omega}$.

Let $Y \coloneqq \tilde{H}(x;\lambda)$ for brevity. First,

$$\text{Var}(Y) = \text{Var}\big(\mathbb{E}[Y \mid \tilde{R}]\big) + \mathbb{E}\big[\text{Var}(Y \mid \tilde{R})\big]. \tag{3.45}$$

Next, for each fixed value of $\tilde{R}$,

$$\text{Var}(Y \mid \tilde{R}) = \text{Var}\big(\mathbb{E}[Y \mid \tilde{R},\tilde{\omega}] \mid \tilde{R}\big) + \mathbb{E}\big[\text{Var}(Y \mid \tilde{R},\tilde{\omega}) \mid \tilde{R}\big]. \tag{3.46}$$

Combining (3.45) and (3.46) yields the decomposition

$$\text{Var}(Y) = V_{\text{tree}} + V_{\text{scen}} + V_{\text{proc}}, \tag{3.47}$$

where

$$V_{\text{tree}} \coloneqq \text{Var}\big(\mathbb{E}[Y \mid \tilde{R}]\big), \tag{3.48}$$

$$V_{\text{scen}} \coloneqq \mathbb{E}\Big[\text{Var}\big(\mathbb{E}[Y \mid \tilde{R},\tilde{\omega}] \mid \tilde{R}\big)\Big], \tag{3.49}$$

$$V_{\text{proc}} \coloneqq \mathbb{E}\Big[\mathbb{E}\big[\text{Var}(Y \mid \tilde{R},\tilde{\omega}) \mid \tilde{R}\big]\Big] = \mathbb{E}_{\tilde{R},\tilde{\omega}}\big[\text{Var}(Y \mid \tilde{R},\tilde{\omega})\big]. \tag{3.50}$$

The term V_{tree} represents structural (tree) uncertainty, V_{scen} captures uncertainty due to the choice of scenario path within each tree, and V_{proc} captures the contribution of process-data uncertainty conditional on a given tree and scenario. The three components in (3.47) sum to the total variance by construction.

Let $\hat{\mu}_r$ denote the empirical mean of Y over all MC iterations with $R^{(\ell)} = r$, and let n_r be the corresponding number of iterations:

$$\hat{\mu}_r = \frac{1}{n_r} \sum_{\ell : R^{(\ell)} = r} \tilde{H}^{(\ell)}(x; \lambda), \qquad r \in \mathcal{R}. \tag{3.51}$$

Using the tree probabilities $\pi_r = 1/R$, an estimator of V_{tree} is

$$\widehat{V}_{\text{tree}} = \sum_{r \in \mathcal{R}} \pi_r \left(\hat{\mu}_r - \hat{\mu}_H(x; \lambda) \right)^2. \tag{3.52}$$

Similarly, for each tree r and each scenario $\omega \in \Omega^{(r)}$, let $\hat{\mu}_{r,\omega}$ denote the empirical mean of Y over all iterations with $R^{(\ell)} = r$ and $\omega^{(\ell)} = \omega$, and let $n_{r,\omega}$ be the corresponding sample size. If $n_{r,\omega} = 0$, the mean can be approximated analytically using (3.38) and empirical process-impact means. The within-tree variance of scenario means is then estimated as

$$\widehat{V}_{\text{scen}}^{(r)} = \sum_{\omega \in \Omega^{(r)}} p_{\omega}^{(r)} \left(\hat{\mu}_{r,\omega} - \hat{\mu}_r \right)^2, \qquad r \in \mathcal{R}, \tag{3.53}$$

and V_{scen} is estimated by averaging over trees:

$$\widehat{V}_{\text{scen}} = \sum_{r \in \mathcal{R}} \pi_r \, \widehat{V}_{\text{scen}}^{(r)}. \tag{3.54}$$

Conditionally on a fixed tree r and scenario ω, the aggregate impact Y is a linear function of the random process impacts, cf. (3.38). Under a first-order approximation that treats the elements of $\tilde{\boldsymbol{\eta}}(x)$ as independent and neglects cross-covariance terms, the conditional variance can be written as

$$\operatorname{Var}(Y \mid \tilde{R} = r, \tilde{\omega} = \omega) \approx \sum_{i=1}^{P} s_{r,\omega,i}(x, \lambda)^2 \, \sigma_i^2, \tag{3.55}$$

where

$$\sigma_i^2 \coloneqq \operatorname{Var}\big(\tilde{\eta}_i(x)\big), \qquad i = 1, \ldots, P, \tag{3.56}$$

are the process-level impact variances, estimated empirically from $\{\tilde{\boldsymbol{\eta}}^{(k)}(x)\}_{k=1}^{K}$.

Define

$$v_{r,\omega} \coloneqq \sum_{i=1}^{P} s_{r,\omega,i}(x, \lambda)^2 \, \sigma_i^2. \tag{3.57}$$

Then an estimator of V_{proc} is given by

$$\widehat{V}_{\text{proc}} = \sum_{r \in \mathcal{R}} \pi_r \sum_{\omega \in \Omega^{(r)}} p_{\omega}^{(r)} \, v_{r,\omega}, \tag{3.58}$$

which corresponds to the expectation in (3.50).

Finally, the three estimators $\widehat{V}_{\text{tree}}$, $\widehat{V}_{\text{scen}}$, and $\widehat{V}_{\text{proc}}$ can be compared to the total variance estimate in (3.44) via

$$\widehat{\operatorname{Var}}\big[\tilde{H}(x; \lambda)\big] \approx \widehat{V}_{\text{tree}} + \widehat{V}_{\text{scen}} + \widehat{V}_{\text{proc}}, \tag{3.59}$$

with the relative magnitudes of the three components interpreted as the shares of structural, scenario-choice, and process-data uncertainty, respectively.

Process-Level Contributions to Variance

Having isolated the component V_{proc} associated with process-data uncertainty, it is natural to further apportion this variance to individual processes. This corresponds to a contribution-to-variance (CTV) analysis in the sense of Heijungs (2024a).

From (3.55) and (3.57), the within-scenario variance can be written as a sum of process-specific contributions. Averaging over trees and scenarios yields

$$V_{\text{proc}} = \sum_{i=1}^{P} \mathbb{E}\big[s_i(x,\lambda)^2\big]\,\sigma_i^2, \tag{3.60}$$

where

$$\mathbb{E}\big[s_i(x,\lambda)^2\big] = \sum_{r\in\mathcal{R}} \pi_r \sum_{\omega\in\Omega^{(r)}} p_\omega^{(r)}\, s_{r,\omega,i}(x,\lambda)^2 \tag{3.61}$$

is the expected squared activity level of process i over the tree-scenario distribution.

The analytic first-order variance contribution of process i is then defined as

$$C_i^{\text{analytic}} := \mathbb{E}\big[s_i(x,\lambda)^2\big]\,\sigma_i^2, \qquad i = 1,\dots,P, \tag{3.62}$$

so that

$$\sum_{i=1}^{P} C_i^{\text{analytic}} = V_{\text{proc}}. \tag{3.63}$$

Each C_i^{analytic} can be interpreted as the portion of the within-scenario variance that is attributable to the uncertainty in the impact factor of process i, assuming mutual independence between processes.

The validity of this decomposition rests on three key assumptions: (i) process-level impacts are mutually independent across processes, (ii) scenario trees are equally likely ($\pi_r = 1/R$), and (iii) the life cycle impact aggregation model is linear. The independence assumption simplifies the analysis by ignoring cross-covariances in the background database but allows for a clear attribution of variance to individual unit processes.

Chapter 4

Case Study: University Campus Furniture

This chapter provides an empirical illustration of the methodological framework developed in Chapter 3. It illustrates how the stochastic SCLCO framework is applied to an institutional procurement decision context characterized by uncertainty in demand, return flows, and life cycle impacts. The case centers on the furniture framework agreement at TU Delft, providing a setting where climate ambitions are translated into operational supply-chain decisions. Relevant code and data are available.[1]

The remainder of this chapter is structured as follows. Section 4.1 motivates the case study. Section 4.2 defines the decision problem, scope, and methodological mapping. Section 4.3 details the product portfolio, circular pathways, and supply chain context. Section 4.4 describes the life cycle inventory parameterization. Section 4.5 presents the historical data analysis and the stochastic modeling of demand and returns. Finally, Section 4.6 details the optimization setup and scenario tree configuration.

4.1 Relevance

This case study starts from the distinctive position of HEIs in sustainability transitions. While universities shape societal practices through education and research, they also operate as large, resource-intensive organizations with extensive infrastructure and diverse activities involving staff, students, and external partners. This dual position allows HEIs not only to contribute to sustainability-oriented knowledge domains but also to implement and evaluate sustainability strategies within their own operations. By embedding research-informed approaches in governance and campus management, universities provide a setting in which climate ambitions can be exemplified, put into practice, and evaluated.

TU Delft has articulated such ambitions through its *Climate Action Program*, aiming to embed sustainability across all activities under the principle *practice what you teach* (van den Dobbelsteen et al., 2023). A central objective is achieving carbon neutrality by 2030, defined as net-zero GHG emissions across scopes 1, 2, and 3 (van Gameren & van den Dobbelsteen, 2023; van den Dobbelsteen & van Gameren, 2022). Although the university has intensified its efforts, current estimates indicate that a substantial share of TU Delft's GHG emissions originate from scope 3 sources, i.e., indirect emissions arising along the supply chain (Herth & Blok, 2022). In particular, construction activities and campus expansion are identified as major contributors. Because many of these impacts are mediated through long-term contractual arrangements, strategic procurement and operational planning constitute key leverage points for reducing emissions while maintaining core operations.

The case offers a suitable analytical setting because it combines several demanding features of sustainable procurement within a single decision context. Specifically, (i) it

[1] https://github.com/philmas/stochastic-supply-chain-life-cycle-optimization

involves a discrete supplier-selection decision under a multi-year framework agreement, implying contractual lock-in; (ii) the decision occurs under uncertainty regarding the level and composition of product and service demand, including return flows; and (iii) the agreement explicitly accommodates circularity options, subject to feasibility and capacity constraints. Together, these features generate environmental and economic trade-offs that cannot be captured through static, product-level comparisons.

At the same time, the sector context balances analytical tractability with methodological insight. Campus furniture involves a relatively limited set of products and services; the high-end furniture sector provides comparatively transparent environmental disclosures; and internal knowledge from a recent framework agreement is available. These conditions allow environmental impacts, economic outcomes, and service terms to be examined jointly, supporting both retrospective analysis and forward-looking assessment.

Accordingly, the furniture framework agreement at TU Delft provides an appropriate empirical setting for evaluating the proposed stochastic SCLCO framework.

4.2 Decision Problem and Scope

The empirical case concerns a multi-year framework agreement for circular campus furniture at TU Delft, awarded to a single supplier. Furniture used in university offices, educational spaces, and shared environments is procured centrally through this agreement, which typically spans five years, with optional extensions. The agreement governs the procurement of furniture products and associated services, including office workstations (e.g., chairs and desks), educational and conference furniture, storage units, and services such as refurbishment and warehousing. It supports more than 30,000 campus users across approximately 550,000 m^2 of net floor area in Delft and The Hague.

The Campus Real Estate & Facility Management (CREFM) office coordinates the agreement on behalf of faculties and departments. Annual expenditure on routine furniture services is estimated at €2.5 million, underscoring the material economic and environmental stakes; total expenditure may fluctuate due to renovation projects. The central decision problem is the *ex-ante* selection of a single supplier (A, B, or C) before future demand levels, return flows, and associated impacts are known.

Upon supplier selection, the university commits to that supplier's product and service portfolio, pricing structures, delivery-mode availability, and circular production capacity limitations for the contract duration. All subsequent decisions are modeled as recourse decisions that scale a predefined set of unit processes (the technology matrix) to satisfy monthly demand and manage return flows, subject to supplier-specific constraints. Circular-share limits, pricing rules, and delivery-mode restrictions enter as hard constraints; no additional policy parameters are decision variables.

Operationally, the decision variables form a scaling vector specifying the operating levels of the unit processes in the technology matrix. Each unit process produces and consumes a certain bundle of flows, such that scaling the activity vector determines the system-wide flow balance. Functional demand is imposed as a constraint on the relevant delivered product–service flow(s) and may be satisfied through alternative technology pathways (e.g., new production, refurbishment, or revitalization), with pathway choice determined endogenously. Meeting demand while minimizing impacts induces activity levels in upstream and enabling processes via the intermediate-flow dependencies encoded in the matrix. The full technology matrix is provided as an Excel worksheet addendum, with flows and processes summarized in Appendix A.

Two modeling conventions apply. First, supplier selection is fixed over the full optimization horizon (i.e., supplier switching is not allowed). Second, the technology matrix is defined at the level of aggregate unit processes rather than continuous intensity measures. The foreground technology matrix is deterministic; all variability enters through scenario-dependent demand and return quantities and through scenario-dependent impact coefficients in the intervention matrix.

The resulting decision problem is formulated as a multi-stage stochastic supplier-

selection problem with recourse.

4.2.1 Methodological Mapping

To ensure traceability between the abstract model and the empirical application, Table 4.1 maps the mathematical notation defined in Chapter 3 to the concrete entities of the campus furniture case.

Table 4.1: Mapping of concepts to the furniture case study.

Concept	Symbol	Case-study instantiation
Supplier set	$\mathcal{X}$	{Supplier A, Supplier B, Supplier C}
Time stage	τ	Starting stage and calendar months (Jan–Jun 2025)
Technology matrix	$\mathbf{A}$	111 unit-process matrix producing 81 flows
Intervention matrix	$\tilde{\mathbf{B}}(x)$	Environmental and economic inventories
Demand vector	$\tilde{\mathbf{f}}_\tau$	Monthly demand and returns by product category
Scaling vector	$\mathbf{s}_\tau$	Monthly scaling of production and service activities
Capacity vector	$\mathbf{u}(x)$	Supplier flow-capacity limits
Trade-off parameter	λ	Weighting of environmental versus economic impact

Note: Notation follows Chapter 3.

4.2.2 Objective Function Criteria

The optimization minimizes a weighted sum of environmental and economic impacts. In this study, environmental impacts are quantified using the IPCC 2021 GWP_{100} characterization method (excluding biogenic CO_2), expressed in $kgCO_2eq$ (Sonderegger, 2024, p. 27). Life cycle impact data are drawn from the Ecoinvent database (Wernet et al., 2016), and assessment is executed using Brightway (Mutel, 2017). Economic impacts represent the total procurement and service costs borne by TU Delft, expressed in Euro (€), over the planning horizon. No further adjustments for inflation or consumer-producer price indices are applied, as the primary interest is supplier differentiation, to which these effect would apply uniformly.

4.2.3 System Boundaries and Exclusions

The case study is based on the furniture products and directly associated services covered by the university's framework agreement for on-campus workplaces, meeting areas, and educational spaces. To maintain tractability and focus on methodological insights, the following elements are explicitly excluded: (i) complete home-working furniture for staff and small workplace accessories; (ii) interior design advice, project management, and furniture management support; (iii) packaging impacts, due to limited disclosure; (iv) avoided-burden credits for recycling; and (v) backlogging of demand. These elements are omitted because they represent different use contexts or would require substantially different modeling approaches.

Within the defined system boundary, the analysis includes furniture items and directly associated services such as delivery, storage, on-site installation, repair, refurbishment, revitalization, and end-of-life handling.

4.3 Product Portfolio and Supply Chain

This case study focuses on nine consolidated furniture categories representing the core campus demand for offices, educational spaces, and shared environments (Table 4.2). The university prioritizes circularity according to the R-ladder, where maintenance and reuse of existing furniture are preferred over new procurement (Kokhuis et al., 2025; van

Gameren & van den Dobbelsteen, 2023). Existing furniture is retained through repair and revitalization; when internal reuse is insufficient, pre-used furniture is sourced externally and refurbished to an as-new state. New furniture is procured only when these options are inadequate, with an emphasis on circular design principles such as modularity and reusability.

The framework agreement includes three suppliers (A, B, and C), which differ in pricing, product portfolios, and environmental characteristics.

Table 4.2: Number of LCA-modeled product variants per product category and supplier.

Product category	Supplier A	Supplier B	Supplier C	Total
Conference chairs	4	4	3	11
Conference tables	1	1	1	3
Education chairs	3	4	3	10
Education tables	3	2	1	6
Electric office desks	4	2	3	9
Manual office desks	1	2	1	4
Office chairs	9	4	2	15
Pedestals	1	1	1	3
Storage cabinets	1	1	1	3
Total	27	21	16	64

4.3.1 Variant Aggregation

Within each product category, multiple variants exist (e.g., wheeled versus sled-based conference chairs) with distinct environmental profiles and prices. Because historical demand data at the variant level are sparse and inconsistent, variant choice is modeled as uniform within each category. This avoids overweighting specific variants based on incomplete records while preserving portfolio diversity. Accordingly, for the products listed in Table 4.2, each variant is assigned equal probability and is sampled and aggregated to the corresponding product-category processes. Their differences (e.g., material composition, quantities, and mass) are still reflected in the life cycle inventory and impact calculations.

4.3.2 Circular Pathways

The model operationalizes the circular hierarchy by distinguishing three acquisition pathways: (i) *revitalization* of existing university-owned items; (ii) *refurbishment* of externally sourced pre-used items; and (iii) *new production*.

Refurbishment involves acquiring pre-used furniture from external sources, restoring it to as-new condition, and supplying it at an as-new price. To the university, refurbished items are functionally equivalent to new products. Environmental impacts arise from component replacement, processing, and additional logistics.

Revitalization refers to the repair and upgrading of furniture already owned by the university. No new product is purchased; only the components and labor required for restoration are accounted for. Environmental impacts are modeled using the same component-level processes as refurbishment, while economic impacts reflect only the replaced components and services (e.g., logistics, technical checks, storage). Revitalization is favored whenever technically feasible and economically or environmentally superior.

For each product category, specific components eligible for replacement under refurbishment or revitalization are defined. These replacement events are modeled probabilistically (Table A.2 and A.3) and embedded in the intervention matrix through MC sampling. Sampling proceeds until the first failed component is realized; failure is guaranteed by construction, so every return triggers either partial repair or full replacement. Product categories for which circular options were not offered or data were unavailable

(e.g., education chairs, pedestals, and storage cabinets) are treated as out of scope for these pathways.

4.3.3 Capacity Context

Optimization requires constraints that reflect the technological, organizational, and contractual limitations of the framework agreement. If demand could always be met without limitation, the problem would reduce to selecting the most favorable technological option and scaling it linearly.

Capacity constraints in this study are derived from the tender documentation, which specifies circularity objectives and performance targets. These are translated into supplier commitments and binding key performance indicators (KPIs), forming the basis for the technological capacity limits imposed in Section 4.6.

Production capacity is defined in technological terms as suppliers' ability to scale operational processes to deliver furniture across the circular hierarchy. This encompasses the capacity to reuse university-owned furniture, refurbish externally sourced items, and supply new furniture when circular options are unavailable. In addition, production capacity includes the technological characteristics of supplier delivery (e.g., combustion-engine versus electric vehicles). Endogenous capacity expansion is not considered within the scope of this analysis.

4.4 Life Cycle Modeling and Parameterization

This section instantiates the life cycle inventory modeling and sampling methodology described in Section 3.4. The final inventory of the case study comprises 111 processes per supplier, totaling 333 processes across nine consolidated furniture categories. Each process is associated with environmental and economic impact categories, resulting in 666 impact coefficients within the characterized intervention matrix $\mathbf{Q}\tilde{\mathbf{B}}(x)$. These processes and their associated impact coefficients are, in turn, parameterized using underlying data sources, including individual bill of materialss (BoMs), pricing information, and event probabilities.

Environmental impacts of activities are linked to ecoinvent 3.11 background data (Wernet et al., 2016) to provide geographical and technological proxies. This approach is methodologically consistent with certified supplier disclosures (e.g., environmental product declarations (EPDs)) and comparable studies, given the unavailability of primary manufacturing data. The assessment of economic impacts is informed by pricing schedules obtained from the internal request for proposal. As these schedules do not provide complete or fully differentiated pricings, assumptions were required, such as treating different office chair variants as having equal unit costs. In addition, certain cost components, including logistics and technical inspections, required supplementary assumptions beyond those explicitly specified in the request for proposal. As actual supplier prices are commercially sensitive and cannot be disclosed or directly attributed, the resulting economic values must be interpreted as derived estimates. Accordingly, the case study reflects market-consistent and procurement-realistic proxy measures that are analytically representative of the underlying procurement context, rather than exact, attributable, or contract-identical prices.

Table 4.3: Illustrative process matrix **P** (flows × processes). The upper block represents the technology matrix **A**, while the lower block shows an illustrative intervention matrix **B**. The representation is product-category agnostic. (R) denotes return flows, which are modeled with negative reference flows. Ellipses indicate additional technologies present in the full matrix but omitted here for clarity.

P matrix	New	Ref.	Rev.	Disp.	Ware.	Coll.	Diesel	Elect.	Inst.	Ret.
Produced	1	1	1			-1				
Collected					1	1			-1	
Warehoused					-1					
Delivery							1	1	-1	
Installed									1	
Returned										-1
Collected (R)			-1	-1	⋯		⋯			1
Warehoused (R)					⋯		⋯			
Delivered (R)							⋯			-1
Methane		⋯		⋯	⋯		⋯		⋯	⋯
Carbon dioxide		⋯		⋯	⋯		⋯		⋯	⋯
⋯		⋯		⋯	⋯		⋯		⋯	⋯
EURO[a]		⋯		⋯	⋯		⋯		⋯	⋯

[a] Multiple timestamped economic indicators (e.g. EURO_{2025}) would generally be preferable, allowing later characterization by **Q**. Since this analysis exclusively covers 2025, such notation is omitted here.

Table 4.4: Presence of non-zero environmental and economic intervention coefficients by life cycle process class. X means non-zero values present.

Process	**Environmental**	**Economic**
New production	X	X
Refurbishment	X	X
Revitalization	X	X
Forward processes		
Collection		
Installation		
Warehousing	X	X
Transportation	X	
Reverse processes		
Collection (R)		X
Return (R)		
Warehousing (R)	X	X
Transportation (R)	X	X
Disposal (R)	X	

4.4.1 Raw Materials and Upstream Logistics

The case study inventory includes material, electronic, chemical, composite, and other production-related inputs (detailed in Table A.5 and A.6). Upstream impacts, including raw material extraction, primary processing, and transportation, are modeled using regionalized average datasets as proxies. Material parameters and upstream transport intensities are summarized in Tables 4.5 and 4.6. Upstream transportation distances follow the modal mixes and 2023 volume updates described by Moreno Ruiz et al. (2018)

and Eurostat statistics (Eurostat, 2022a, 2022b, 2023). Detailed material-specific values and implementation data are provided in Appendix A. Explicit economic impacts are not assigned to raw materials, as raw material costs, labor, processing and distribution are assumed to be embedded within the selling prices of the products. The resulting foreground processes described here are not included in the process matrix $\mathbf{P}$ but instead aggregated into the manufacturing processes of products.

Table 4.5: Overview of material parameters applied in the life cycle inventory. Material-specific values and data sources are reported in Table A.6.

Parameter	Description
Recycled content	Share of secondary material used as input in the production of a given material.
PCR/PIR content	Share of post-consumer and post-industrial recycled material assumed in polymer production.
Recycling rate	Fraction of a material that is assumed to be collected and recycled at end of life.
Density	Bulk material density used to convert between mass- and volume-based inventories.
Glass fiber content	Mass fraction of glass fiber reinforcement assumed in fiber-reinforced polymer materials.
Talc/glass fiber content	Mass fraction of mineral or glass fiber fillers assumed in polypropylene-based materials.
Renewable electricity share	Share of renewable electricity assumed in the energy supply of selected production processes.

Table 4.6: Transport intensities by material class and transport mode. Values represent tonne–kilometres (tkm) associated with each material class.

Material class	Rail (tkm)	Road (tkm)	Inland water (tkm)
Natural fibers	0.0011649	0.2862246	0.0013785
Minerals	0.0042353	0.0860698	0.0011252
Polymers	0.0305550	0.1817282	0.0244955
Chemicals	0.0305550	0.1817282	0.0244955
Wood	0.0240677	0.1972218	0.0020459
Metals	0.0463444	0.1781996	0.0106329
Electronics	0.0028763	0.1976555	0.0026011

4.4.2 Manufacturing and End-of-Life

Foreground unit processes are constructed by mapping supplier-provided BoMs to the corresponding raw material inventories. For each product category, a dedicated unit process is defined and included in the process matrix $\mathbf{P}$. The environmental impacts for these processes are determined by the raw material exchange coefficients reflecting upstream composition and raw material profiles. The economic impacts are determined by the supplier information and for some processes additionally conditioned by the necessary replacement of components.

Foreground processes are normalized to complete functional units, such that one unit of process activity yields one unit of product output. The first three columns of Table 4.3 illustrate the manufacturing technologies of the matrix $\mathbf{A}$ for an arbitrary circular-eligible product category. Uncertainty in background inventory data, typically represented using pedigree matrices in the ecoinvent database, is propagated via MC sampling and reflected in the intervention matrix $\tilde{\mathbf{B}}(x)$. For new furniture production, environmental impacts are

calculated directly from the assembled foreground inventory and subsequently combined with supplier-specific economic impact data.

End-of-life treatment follows a cut-off approach, where recycling rates (Table 4.5 and A.6) determine the fraction of material diverted from incineration or landfill. Because furniture products have long service lives relative to the planning horizon, raw material production and disposal are modeled as distinct processes occurring at different time instances. For refurbishment and revitalization pathways, component-level processes are modeled according to the replacement events defined in Section 4.3.2, including the disposal of replaced components. Returned products that are not revitalized must either be stored for future processing or disposed of; these implicitly modeled flows are indicated schematically in the supply chain representation. To enforce this behavior within the optimization model, returned products are represented via negative reference flows, together with a terminal balance constraint that ensures all such flows are fully resolved by the end of the planning horizon, i.e., no residual unprocessed flows remain.

For both manufacturing and end-of-life handling, the economic dimension is structured as a process-based cost architecture aligned with the product life cycle. For circular production (i.e., refurbishment and revitalization processes), costs are based on the occurrence of component-specific failure and replacement events. End-of-life disposal costs are not incurred explicitly and are instead assumed to be embedded within the logistics and distribution cost structure.

To illustrate how uncertainty in foreground inventories propagates to process-level impacts, Figure 4.1 presents an illustration of impact distributions for revitalization of an office chair. The distributions arise from MC sampling of background inventory uncertainties embedded in $\tilde{\mathbf{B}}(x)$.

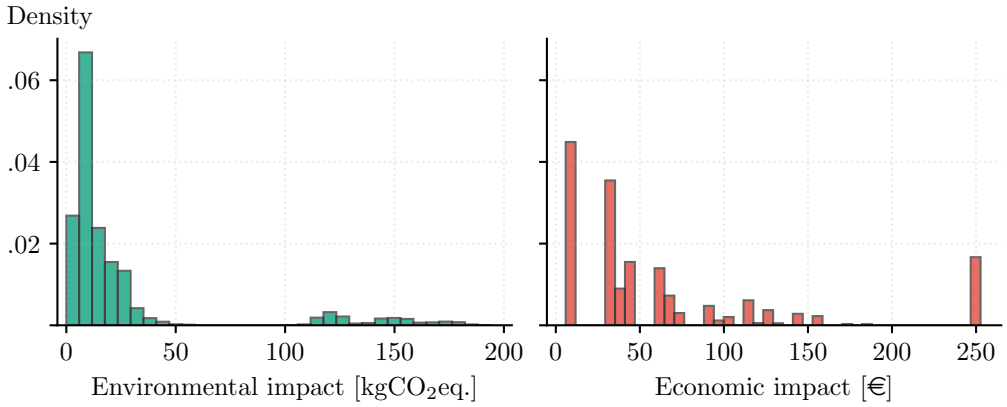

Figure 4.1: Impact distributions for the office chair revitalization manufacturing process obtained via MC propagation of inventory uncertainty. The foreground unit process corresponds to the revitalization technology (third column in Table 4.3).

4.4.3 Warehousing and Transportation

Distribution and storage stages represent the foreground logistics between suppliers and the campus. Transportation is modeled using diesel (EURO 6) and battery-electric vehicles (BEVs), for which a custom BEV dataset was constructed by modifying the EURO 6 unit process. Key operational parameters, including payload, vehicle lifetime, and lifetime throughput used to amortize battery manufacturing impacts, are summarized in Table A.8.

For the BEV technology, the use phase is replaced by electricity consumption from the Dutch medium-voltage grid plus a battery-manufacturing proxy, with brake-wear exchanges scaled to reflect regenerative braking. Operational intensities and allocation factors are reported in Table A.8. The electricity intensity is deliberately conservative to keep results robust under plausible efficiency improvements (Piepenbrink et al., 2025;

Volvo Trucks, 2022). Supplier-to-campus transport is translated into per-unit impacts using the mass of flows and road distances between each supplier and the campus.

Warehousing is modeled as a separate technosphere process defined per m^2 month of occupied floor area. The process includes amortized building infrastructure, electricity use, and heat demand, with energy intensities and operational assumptions specified in Table A.9. Storage impacts are treated independently of the operator, as identical ambient conditions and energy requirements are applied. Product-specific warehousing impacts are derived by scaling the per–m^2 month process according to the area densities specified in Table A.9.

The economic impact for forward logistics is zero as this is already embedded in the pricing of the furniture products. Conversely, reverse logistics costs are modeled explicitly, assuming a uniform transportation cost per unit across all furniture categories. Moreover, for the reverse logistics, a technical check at collection is required this is also included.

4.5 Stochastic Demand and Return Analysis

The final demand modeled in the case study comprises both forward demand from the university campus and reverse flows associated with returned furniture. The primary objective is to generate future demand and return scenarios that are representative of historical purchasing behavior, including tails and intermittency, and suitable for multistage stochastic optimization.

Demand is modeled based on reconstructed historical procurement data. Individual purchase records are classified into the nine furniture product categories defined in Section 4.3. These classified time series form the empirical basis for subsequent statistical analysis.

Capacity planning and procurement decisions under uncertainty require probabilistic forecasts that capture both the temporal patterns of furniture orders and the variability in order sizes. Traditional point forecasting is insufficient here, as it fails to represent the range of plausible futures needed for robust decision-making. Accordingly, this study develops stochastic demand and return forecasts enabling evaluation of supplier selection, capacity allocation, and circular procurement strategies.

The analysis proceeds by characterizing historical order timing and quantities, fitting appropriate statistical distributions, and generating MC scenarios that preserve empirical correlation structures. Product returns are forecasted using cohort-based reliability models based on Weibull lifetime distributions.

The forecasting exercise produces a 12-month horizon (January–December 2025) at monthly resolution. Each scenario represents a complete realization of (i) monthly demand quantities per product category, (ii) monthly return quantities per product category, and (iii) cross-product dependence structures. These scenario realizations constitute the stochastic input for scenario tree generation (Section 4.6).

4.5.1 Historical Demand Analysis

Historical item records (2015–2023) were used to construct a dataset of 15,532 entries. After classification and cleaning, 4,280 valid transactions remained (27.6% of the original dataset). Table 4.7 reports data completeness by consolidated product category. Sample sizes range from 262 (Pedestals) to 829 (Office Chairs), sufficient for intermittent-demand modeling.

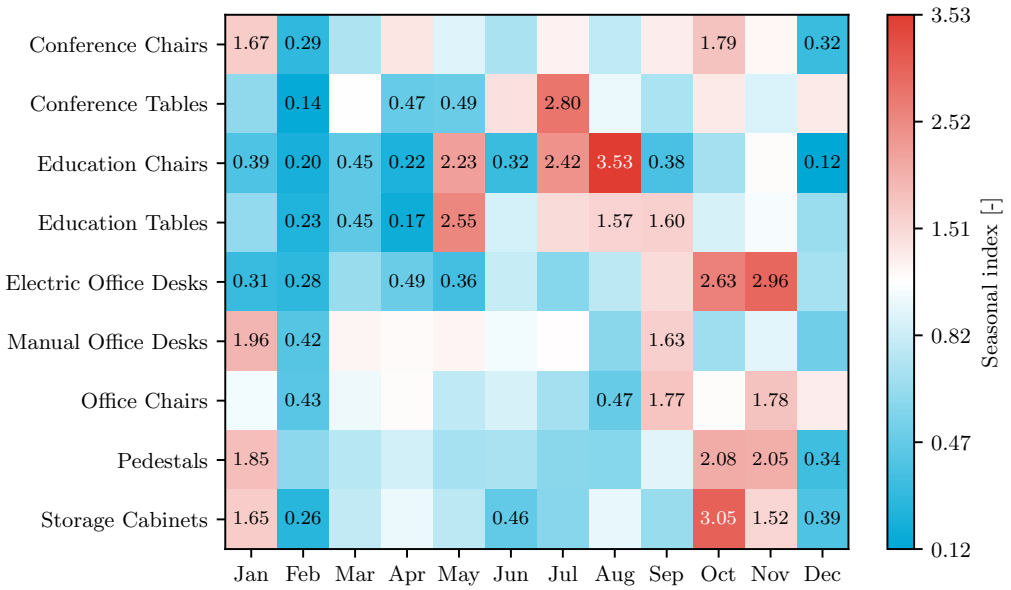

Figure 4.2: Seasonal indices of monthly demand by furniture product category, computed as the ratio of month-specific demand to the long-run monthly average.

Table 4.7: Completeness of key date fields by furniture product category. Percentages indicate the share of records with non-missing values.

Product category	Records	Issue date (%)	Order date (%)	Delivery date (%)
Conference chairs	595	100.0	100.0	89.9
Conference tables	286	100.0	100.0	86.7
Education chairs	336	100.0	100.0	74.1
Education tables	643	100.0	100.0	89.9
Electric office desks	399	100.0	100.0	93.2
Manual office desks	597	100.0	100.0	87.9
Office chairs	829	100.0	100.0	91.4
Pedestals	262	100.0	100.0	89.7
Storage cabinets	350	100.0	100.0	90.6

Demand is first transformed from a daily periodicity to a weekly time grid comprising 468 periods. Individual product codes are subsequently aggregated into nine consolidated product categories to ensure a sufficient number of non-zero observations while preserving functional distinctions. As a result, the data consist of weekly order quantities for each product category.

This aggregation occurred in order to (i) grouping jointly planned codes (e.g., conference chair variants); (ii) grouping codes with correlated temporal patterns; and (iii) combining small-sample codes to increase robustness.

Weekly demand observations are aggregated to the month-of-year level to examine temporal patterns. Seasonal indices, computed as the ratio of monthly demand to the long-run average, are summarized in Figure 4.2 and Table 4.8. This involves pooling all historical observations for January, all for February, and so forth, to identify seasonal factors.

Three recurring patterns emerge. Office furniture displays moderate seasonality, peaking in Q4 and January. Education furniture shows strong seasonality aligned with the academic calendar (summer peaks). Storage and conference furniture exhibits intermediate to strong seasonality, concentrated in Q4.

Table 4.8: Seasonality characteristics by furniture product category. Peak and low months indicate the highest and lowest relative demand; indices report the minimum and maximum seasonal index; CV is the coefficient of variation.

Product category	Peak month	Low month	Min index	Max index	CV (%)
Conference chairs	October	February	0.29	1.79	47.4
Conference tables	July	February	0.14	2.80	67.8
Education chairs	August	December	0.12	3.53	111.1
Education tables	May	April	0.17	2.55	68.8
Electric office desks	November	February	0.28	2.96	89.8
Manual office desks	January	February	0.42	1.96	44.9
Office chairs	November	February	0.43	1.78	43.0
Pedestals	October	December	0.34	2.08	61.6
Storage cabinets	October	February	0.26	3.05	76.9

4.5.2 Demand Characteristics

Order sizes are characterized by strong right skewness and tails (Table 4.9). Medians are substantially below means, and skewness often exceeds 2, indicating infrequent but large orders. High excess kurtosis confirms lumpy, project-driven procurement. These diagnostics motivate using tailed parametric distributions (e.g., lognormal, gamma) in the demand-size model.

Table 4.9: Distribution statistics of order quantities by furniture product category.

Product category	Mean	Median	Std. dev.	Min	Max	Skew.	Exc. kurt.
Conference chairs	12.81	6	20.07	1	250	6.39	65.67
Conference tables	2.92	1	5.19	1	48	5.01	30.48
Education chairs	32.56	15	54.25	1	530	4.26	26.14
Education tables	7.12	3	17.78	1	320	12.08	189.46
Electric office desks	5.19	2	12.26	1	163	8.19	90.06
Manual office desks	11.34	6	13.67	1	90	2.05	4.80
Office chairs	15.01	6	24.79	1	200	3.31	14.35
Pedestals	12.85	5	23.06	1	176	4.22	20.97
Storage cabinets	6.24	2	16.00	1	236	9.62	125.10

Intermittency is evaluated using the proportion of zero-demand weeks, the ADI, and the Syntetos–Boylan classification (Table 4.10). All demand series are categorized as lumpy, with zero-demand weeks ranging from 30% to 75% and high CV^2 values. This supports the choice of an intermittent demand model (occurrence process plus size distribution).

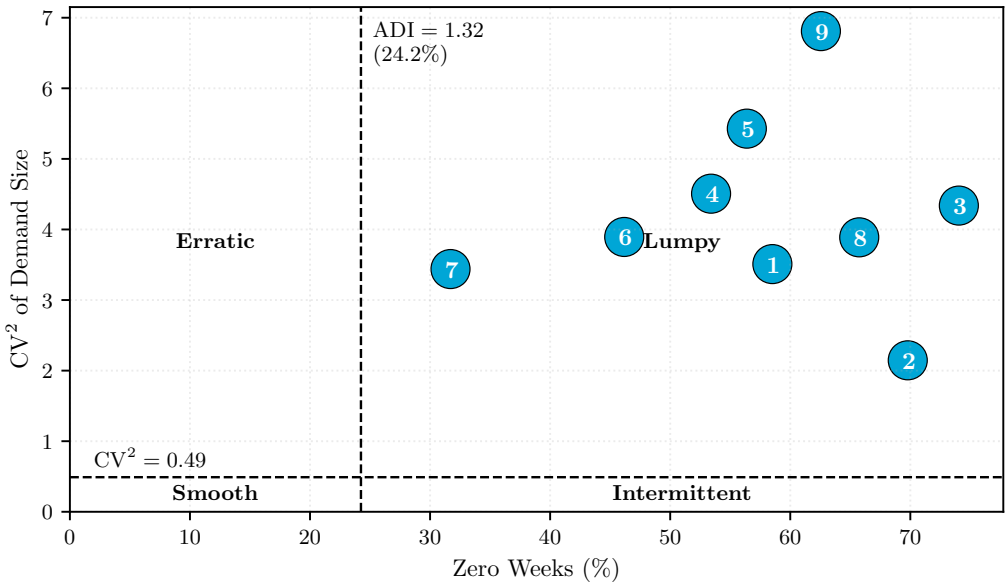

1: Conference Chairs
2: Conference Tables
3: Education Chairs
4: Education Tables
5: Electric Office Desks
6: Manual Office Desks
7: Office Chairs
8: Pedestals
9: Storage Cabinets

Figure 4.3: Intermittency classification of weekly demand series for the nine furniture product categories (Syntetos–Boylan scheme).

Table 4.10: Intermittency characteristics of weekly order patterns by furniture product category.

Product category	Zero weeks (%)	ADI (weeks)	Orders/year	CV^2	Class
Conference chairs	58.5	2.40	21.67	3.51	Lumpy
Conference tables	69.8	3.30	15.78	2.15	Lumpy
Education chairs	74.0	3.78	13.56	4.34	Lumpy
Education tables	53.4	2.13	24.33	4.51	Lumpy
Electric office desks	56.4	2.24	22.78	5.43	Lumpy
Manual office desks	46.2	1.86	28.11	3.89	Lumpy
Office chairs	31.7	1.47	35.67	3.44	Lumpy
Pedestals	65.7	2.93	17.89	3.89	Lumpy
Storage cabinets	62.6	2.67	19.56	6.81	Lumpy

Cross-product interactions are examined using the correlation matrix of monthly demand totals (Figure 4.4). Three clusters emerge: (i) office furniture; (ii) education furniture; and (iii) storage/pedestals. While mean off-diagonal correlations are modest, strong positive dependence exists within clusters, necessitating correlation-preserving scenario generation.

Dependence between demand timing and order size was tested (Table 4.11). Correlations between order size and preceding/succeeding zero-demand gaps are generally weak ($|\rho| < 0.2$), supporting the assumption that occurrence and size can be modeled as approximately independent components.

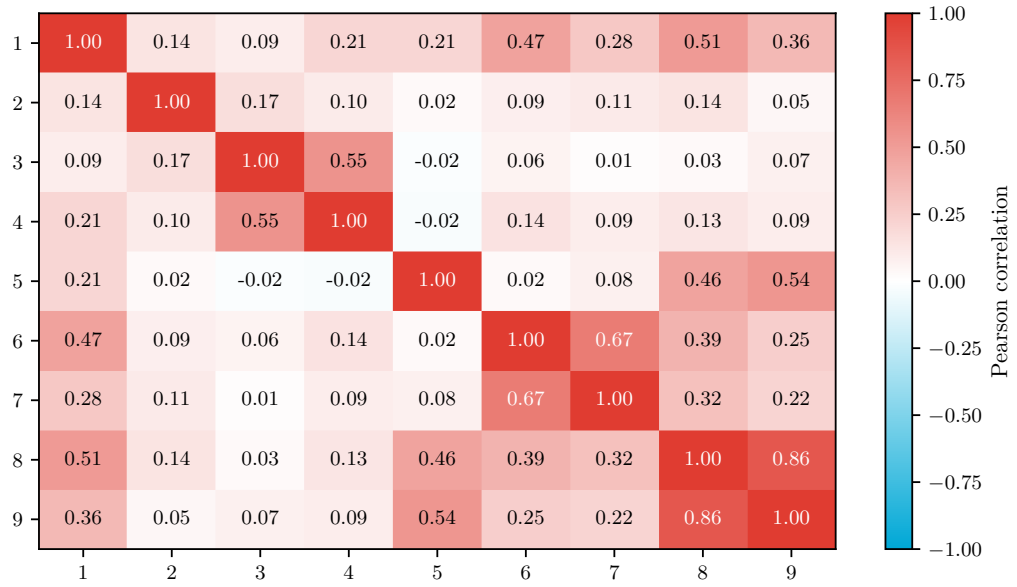

Figure 4.4: Heatmap of the correlation matrix of historical monthly demand totals across the nine furniture product categories.

Table 4.11: Dependence diagnostics between weekly gaps and order size.

Product category	ρ_b	$\rho_b^{(\log)}$	ρ_a	$\rho_a^{(\log)}$	Slope	p	R^2	n
Conference chairs	0.009	-0.129	0.182	0.052	-0.073	0.073	0.017	194
Conference tables	0.044	0.011	-0.068	-0.131	0.003	0.895	0.000	141
Education chairs	-0.116	-0.273	-0.003	-0.061	-0.080	0.002	0.074	121
Education tables	-0.069	-0.116	-0.102	-0.133	-0.076	0.086	0.014	218
Electric office desks	-0.119	-0.180	-0.087	-0.196	-0.083	0.010	0.033	204
Manual office desks	-0.035	-0.060	-0.077	-0.090	-0.049	0.341	0.004	252
Office chairs	-0.047	-0.091	-0.008	-0.013	-0.126	0.106	0.008	320
Pedestals	0.067	0.126	0.115	0.074	0.056	0.113	0.016	160
Storage cabinets	0.052	0.002	-0.062	-0.105	0.001	0.978	0.000	175

Residual temporal dependence in monthly demand (after removing trend and seasonality) was assessed using Ljung–Box tests (Table 4.12). For most categories, no significant autocorrelation remains. Exceptions include electric and manual office desks, suggesting some clustering. To maintain tractability, monthly demand is simulated as conditionally independent, which is a conservative assumption for capacity planning as it increases dispersion.

Table 4.12: Ljung–Box test p-values for residual monthly demand after removing trend and seasonality.

Product category	Lag 1 month	Lag 6 months	Lag 12 months
Conference chairs	0.110	0.058	0.046
Conference tables	0.208	0.575	0.441
Education chairs	0.112	0.210	0.057
Education tables	0.540	0.319	0.320
Electric office desks	< 0.001	< 0.001	0.003
Manual office desks	0.014	0.038	0.193
Office chairs	0.909	0.572	0.267
Pedestals	0.646	0.365	0.041
Storage cabinets	0.835	0.622	0.415

4.5.3 Installed Base and Return Potential

Return volumes are not inferred from historical return records due to data limitations. Instead, returns are generated using a cohort-based reliability approach. Historical sales are aggregated into monthly cohorts (Table 4.13), and returns are modeled as end-of-use events governed by Weibull distributions.

Table 4.13: Active cohort summary (history used for return forecasting).

Product	First	Last	Cohorts	Units sold	Mean size
Office chairs	2014–12	2024–12	116	12,780	110.17
Manual office desks	2014–12	2024–11	107	6,941	64.87
Electric office desks	2015–03	2024–12	97	2,079	21.43
Education chairs	2014–12	2024–11	68	11,431	168.10
Education tables	2014–12	2024–12	99	4,639	46.86
Conference chairs	2015–01	2024–03	93	7,639	82.14
Conference tables	2014–12	2024–11	80	840	10.50
Pedestals	2015–01	2024–11	87	3,391	38.98
Storage cabinets	2014–12	2024–11	91	2,270	24.95

4.5.4 Intermittent Demand Model

Future demand scenarios are generated using a two-part model: (i) an occurrence process describing the timing of orders, and (ii) a size process describing the quantity conditional on occurrence. This approach captures both the frequency of zero-demand periods and the heavy tails of order sizes.

Model estimation uses weekly data to reduce daily variance, while scenarios are reported monthly to align with decision-making timescales. Inter-arrival times are modeled using exponential, geometric, or Weibull distributions, selected via Akaike information criterion (AIC). Positive weekly demand volumes are modeled using lognormal, gamma, or zero-truncated negative binomial distributions, also selected via AIC and stratified by calendar month to capture seasonality.

Cross-product dependence is characterized via the empirical correlation matrix. Scenarios are generated by simulating inter-arrival times and sampling demand quantities, then aggregating to monthly totals. The number of scenarios is chosen to ensure stability of tail quantiles.

Evaluation takes place of the *marginal* demand scenario generator before any cross-product dependence. For each product type, monthly demand over 2015–2023 is represented by the two-component specification described above, combining an occurrence process with a conditional size distribution.

Based on the fitted parameters, $S = 10{,}000$ independent demand scenarios are generated for each month in the planning horizon 2025–01 to 2025–06, yielding simulated draws for any month of the year. These simulated month-ahead distributions are assessed against historical monthly demand from 2015–2023 using a month-of-year (MOY) stationarity assumption: for each calendar month (January–June), the simulated distribution is compared with the empirical distribution of historical observations from the same month-of-year.

Table 4.14 summarizes two aspects of intermediate performance. First, overall agreement in level and dispersion is reported using pooled simulated draws across 2025–H1. Second, MOY-conditional diagnostics are reported to assess seasonal alignment and tail behavior. Seasonal alignment is quantified by (i) the mean absolute deviation between simulated and historical MOY means (MOY MAE%), which measures the typical magnitude of seasonal discrepancies, and (ii) the signed deviation (MOY Bias%), which indicates systematic over- or underestimation across months. Tail behavior is checked using exceedance rates, i.e., the fraction of historical January–June observations lying above the simulated MOY-specific 95th percentile (and below the 5th percentile). For a well-calibrated generator, both exceedance rates are expected to be close to 0.05.

Table 4.14: Validation of the marginal demand scenario generator for 2025–01 to 2025–06 against 2015–2023 month-of-year benchmarks (January–June). The table reports simulated-to-historical ratios for mean, standard deviation, and 95th percentile, together with MOY seasonal mismatch (MOY MAE% and MOY Bias%) and tail exceedance rates relative to simulated MOY-specific 95th and 5th percentiles (target ≈ 0.05).

Product type	Hist. mean	Sim. mean	Mean ratio	Std. ratio	Q95 ratio
Conference chairs	65.59	63.71	0.97	0.86	0.93
Conference tables	4.39	5.81	1.32	1.06	1.27
Education chairs	61.67	60.06	0.97	1.20	1.04
Education tables	32.94	30.38	0.92	0.86	0.97
Electric office desks	9.20	10.90	1.18	0.91	1.08
Manual office desks	69.54	63.01	0.91	0.73	0.89
Office chairs	91.07	84.67	0.93	0.74	0.93
Pedestals	27.74	27.44	0.99	0.81	1.00
Storage cabinets	16.31	17.78	1.09	1.03	1.25

Product type	MOY MAE (%)	MOY bias (%)	$P(\text{obs} > q_{95})$	$P(\text{obs} < q_{05})$
Conference chairs	23.0	-2.9	0.093	0.000
Conference tables	42.6	32.3	0.074	0.000
Education chairs	45.9	-2.6	0.074	0.000
Education tables	14.7	-7.8	0.037	0.000
Electric office desks	22.7	18.4	0.093	0.000
Manual office desks	21.4	-9.4	0.056	0.056
Office chairs	15.2	-7.0	0.074	0.056
Pedestals	25.9	-1.1	0.056	0.000
Storage cabinets	14.0	9.0	0.074	0.000

For the forward-looking scenario set used in the optimization, the marginal diagnostics in Table 4.14 indicate that average demand levels are broadly reproduced for several product types, while non-trivial mean shifts remain for selected categories (e.g., conference tables with a mean ratio of 1.32 and manual office desks with a mean ratio of 0.91; ratios above/below one indicate over-/underestimation relative to history). Demand dispersion is typically understated, as reflected by standard-deviation ratios below one for most product types. Seasonal alignment is also limited for a subset of categories, with MOY MAE exceeding 40% in the most mismatched cases. These intermediary diagnostics emphasize that the above must be seen as a pragmatic simulation, and motivate a sensitivity analysis.

4.5.5 Return Generation Model

The return process is modeled using a Weibull distribution to capture the characteristic wear-out phase of durable furniture assets. For the categories, a shape parameter $\beta > 1$ is selected, reflecting an increasing hazard rate as products age. The absence of historical unit-level failure data, make that the scale parameters η are calibrated to match an estimated service lifetime. Returns are then forecasted using the cohort-based reliability framework parameterized by Weibull distributions (Table 4.15). The resulting scenarios yield the aggregate uncertainty distributions summarized in Table 4.16.

Table 4.15: Weibull parameter specifications used for return forecasting (base case).

Product type	β	η (months)	Characteristic life (years)
Office chairs	3.2	156	13.0
Manual office desks	3.4	204	17.0
Electric office desks	2.8	156	13.0
Education chairs	3.0	144	12.0
Education tables	3.3	168	14.0
Conference chairs	3.6	216	18.0
Conference tables	3.8	240	20.0
Pedestals	3.4	192	16.0
Storage cabinets	3.8	240	20.0

Table 4.16: Annual return uncertainty summaries (scenario distribution, 2025 totals).

Product type	Mean	Std. dev.	P5	P25	P50	P75	P95	CV (%)
Office chairs	563.0	22.9	526	547	563	578	601	4.1
Manual office desks	145.9	11.9	127	138	146	154	166	8.1
Electric office desks	76.2	8.5	62	70	76	82	91	11.2
Education chairs	785.2	26.9	741	767	785	803	830	3.4
Education tables	179.4	13.0	158	171	179	188	201	7.3
Conference chairs	96.9	9.7	81	90	97	103	113	10.0
Conference tables	7.3	2.7	3	5	7	9	12	36.7
Pedestals	82.4	8.9	68	76	82	88	97	10.8
Storage cabinets	23.8	4.8	16	20	24	27	32	20.3

4.6 Optimization Setup

4.6.1 Scenario Tree Configuration

Multistage uncertainty is represented using a joint scenario tree that combines (i) uncertain supplier-dependent LCA coefficients and (ii) uncertain demand/return trajectories. The case study uses a six-month planning horizon with $T = 7$ stages (including the initial decision stage). A node-based tree representation is required to enforce NACs in the multistage stochastic program. The baseline branching vector is

$$[1, 6, 3, 3, 3, 3, 3],$$

which yields $6 \cdot 3^5 = 1{,}458$ leaf scenarios per tree.

The joint tree combines two sources of uncertainty that are modeled as independent. The first-stage decision is a here-and-now supplier selection, taken before any uncertainty is revealed and held fixed over the horizon. The second stage captures uncertainty in supplier-specific impact coefficients (environmental and economic), which are interpreted as portfolio-level realizations affecting all subsequent operational activities. Demand and

return uncertainty enters over the operational stages and is represented by branching in the demand/return process after the initial root values.

Under the independence assumption, joint scenarios are formed by taking the Cartesian product of supplier-impact realizations and demand/return paths. The probability of a joint leaf scenario $\omega \in \Omega$ factorizes as

$$\mathbb{P}(\omega) \;=\; \mathbb{P}(\omega^{\text{impact}})\,\mathbb{P}(\omega^{\text{dem/ret}}) \in [0,1],$$

so that each supplier is evaluated against the same set of demand/return trajectories. This construction supports comparability across suppliers: differences in objective values arise from supplier-dependent coefficients and decision responses, rather than from differences in the realized demand paths.

The choice of tree structure reflects a trade-off between representational fidelity and computational tractability. Increasing branching improves coverage of heavy-tailed and high-variance uncertainty, while increasing depth improves the model's ability to represent staged information revelation and hedging over time. The adopted branching vector places more width in early stages to capture initial supplier portfolio-impact variability while keeping the total number of leaf scenarios manageable. To assess robustness with respect to sampling and discretization error, multiple trees are generated using different random seeds and analyzed comparatively.

4.6.2 Contractual Capacity Constraints

The capacity vector $\mathbf{u}(x)$ is populated by applying suppliers' circularity-share commitments to the 90th percentile of historical demand, which serves as a heuristic for a 'high-stress' operational level. The resulting technological capacity limits (Table 4.17) represent binding commitments and are implemented as hard upper bounds on monthly circular production (revitalization/refurbishment) and electric delivery capacity. When these limits are exhausted in high-demand scenarios, the optimization model is forced to rely on alternative pathways (e.g., new production, diesel transport or warehousing).

Table 4.17: Supplier-specific monthly capacity limits for circular production and electric delivery by furniture product category. Capacities are derived from tender commitments and applied as hard upper bounds in the optimization model.

Product category	Capacity type	Sup. A	Sup. B	Sup. C
Office chairs	Circular production	143	163	173
Office chairs	Electric delivery	122	204	204
Manual office desks	Circular production	92	105	112
Manual office desks	Electric delivery	79	132	132
Electric office desks	Circular production	33	38	40
Electric office desks	Electric delivery	28	47	47
Conference chairs	Circular production	130	149	158
Conference chairs	Electric delivery	112	186	186
Conference tables	Circular production	21	24	25
Conference tables	Electric delivery	18	30	30
Education chairs	Electric delivery	165	276	276
Education tables	Circular production	76	87	93
Education tables	Electric delivery	65	109	109
Pedestals	Electric delivery	51	85	85
Storage cabinets	Electric delivery	26	44	44

4.6.3 Implementation and Solver

This section summarizes the computational implementation and execution setup used to solve the integrated stochastic optimization model. The main goals of the implementation are (i) a clean separation between model specification, scenario generation, and algorithmic orchestration, (ii) scalable parallel execution over large scenario sets, and (iii) reproducibility through deterministic experiment configuration and structured logging. The model is implemented in Python using Pyomo for algebraic modeling (Bynum et al., 2021; Hart et al., 2011), solved with Gurobi (Gurobi Optimization, LLC, 2025), and parallelized using the `mpisppy` framework (Knueven et al., 2023), on an HPC cluster (Delft High Performance Computing Centre [DHPC], 2024).

The implementation follows a modular architecture. One component constructs scenario data (samples of uncertain LCA impacts and demand/return trajectories), a second component instantiates the Pyomo optimization model given a scenario (or a bundle of scenarios), and a third component provides algorithm drivers for running PH, EF solves, and statistical evaluation procedures. This separation makes it possible to reuse the same scenario generation and model definition across different computational experiments (e.g., optimization runs, EVPI/VSS calculations, and gap estimation), while keeping the numerical settings and execution details localized to the run configuration.

A direct extensive-form representation of the full scenario tree is large, so the implementation uses scenario bundling: multiple scenarios are aggregated into a single EF subproblem that is treated as one composite unit within PH. The scenario tree used in the case study has branching factors $[6, 3, 3, 3, 3, 3]$ (one LCA stage followed by five time stages), yielding $6 \cdot 3^5 = 1458$ leaf scenarios. These are partitioned into bundles of 27 scenarios each, giving 54 bundles. Each bundle is compiled into an EF model that retains the multistage structure internally, but is solved as a single subproblem in the PH loop. Bundles are pre-generated and serialized once, which decouples scenario construction from optimization and improves reproducibility across repeated runs.

PH is executed using the hub–spoke architecture provided by `mpisppy` (Knueven et al., 2023). The hub maintains the consensus values of the nonanticipative variables ($\hat{x}$) and updates multipliers and proximal penalties. Spokes provide auxiliary computations to accelerate convergence and to obtain bounds. Two spoke types are used: a Lagrangian spoke that produces dual information and lower-bound estimates, and a feasibility-oriented spoke that generates candidate primal solutions by fixing nonanticipative variables to hub values. Iterations are coordinated by `mpisppy`'s internal scheduler, which synchronizes communication and convergence checks across MPI ranks.

Each PH iteration requires solving one EF bundle subproblem per MPI rank. These subproblems are solved with Gurobi through Pyomo. To avoid oversubscription on shared nodes, solver parallelism is restricted to a single thread per rank, so that parallelism is expressed primarily across bundles rather than within a single MILP solve. Warm-starting is used where supported by the solver interface to reduce repeated solve time across PH iterations.

All experiments are executed under the Slurm workload manager with MPI-based parallelization. Bundle generation is treated as an embarrassingly parallel precomputation and is run as a job array over random seeds. Optimization and evaluation runs allocate one MPI rank per bundle (54 ranks), so each worker solves one EF subproblem per iteration and exchanges the required consensus and dual information with the hub. Postprocessing steps aggregate distributed outputs into analysis-ready datasets.

Table 4.18: Computational setup and execution settings.

Component	Parameter	Value
Modeling	Framework	Pyomo (algebraic modeling)
	Formulation	Scenario-based stochastic program
Algorithm	Decomposition	PH (hub–spoke)
	Subproblem type	Extensive form (EF) per bundle
Scenarios	Tree branching	$[6, 3, 3, 3, 3, 3]$
	Leaf scenarios	1458
	Bundling rule	27 scenarios per bundle $\Rightarrow$ 54 bundles
Solver	MILP solver	Gurobi
	Threads	1 thread per MPI rank
Parallelization	Runtime system	MPI
	PH ranks	1 rank per bundle (54 ranks total)

Chapter 5

Results

This chapter presents the numerical outcomes of the optimization and evaluates the robustness of the resulting policies. It begins by validating the scenario-tree approximation (Section 5.1) and characterizing the environmental–economic trade-offs that shape supplier selection (Section 5.2). Statistical assessments of solution quality are provided through optimality-gap analysis (Section 5.3) and value-of-information metrics (Section 5.4). Finally, Section 5.5 and Section 5.5.2 characterize the distribution of realized impacts and apportion the observed variance to its underlying structural and epistemic sources.

5.1 Scenario Tree Approximation

This section evaluates how faithfully the six-point scenario trees approximate the underlying empirical outcome distributions of the supplier-specific life cycle impact coefficients. The resulting approximation diagnostics provide case-study–specific context for interpreting the subsequent optimization results: any case-study conclusions are necessarily conditioned on the quality of such representation.

5.1.1 Life Cycle Impacts

Figure 5.1 evaluates the accuracy of the six-point scenario reduction (see Section 4.6.1) for supplier-dependent life cycle impact coefficients. For each coefficient, $W_{1,\text{realized}}$ denotes the 1-Wasserstein distance between the full empirical outcome distribution and its corresponding six-point scenario-tree approximation. To contextualize this value, $W_{1,\text{typical}}$ is defined as the median W_1 obtained from repeated stratified six-point resampling (one draw per environmental quantile bin with accompanying economic impact) under the same binning rule on the same underlying coefficient distribution. The realized distances align closely with this typical baseline. Environmental coefficients exhibit tighter concentration around the identity line, whereas economic coefficients show a wider dispersion and several upward outliers.

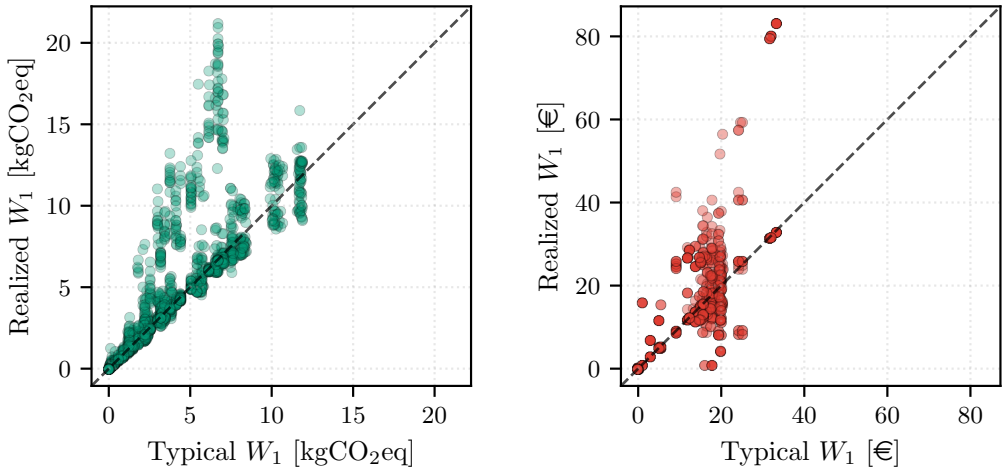

Figure 5.1: Realized versus typical 1-Wasserstein distance W_1 for the scenario-tree reduction of life cycle impact coefficients. Each point corresponds to one process–impact coefficient under one scenario-tree realization (seed). The dashed line indicates $W_{1,\text{realized}} = W_{1,\text{typical}}$: points above (below) the line correspond to worse-than-typical (better-than-typical) six-point approximations. Left: environmental impacts (kgCO$_2$eq). Right: economic impacts (€).

To obtain a scale-free assessment across coefficients with different units and dispersion, Figure 5.2 reports the range-normalized distance $W_1/(\max - \min)$, where $(\max - \min)$ is the observed range of the full coefficient distribution. This can be interpreted as an average transport displacement expressed as a fraction of each coefficient's outcome span. Normalized distances are typically small for environmental coefficients and broadly similar across suppliers. Economic coefficients exhibit larger dispersion and heavier upper tails.

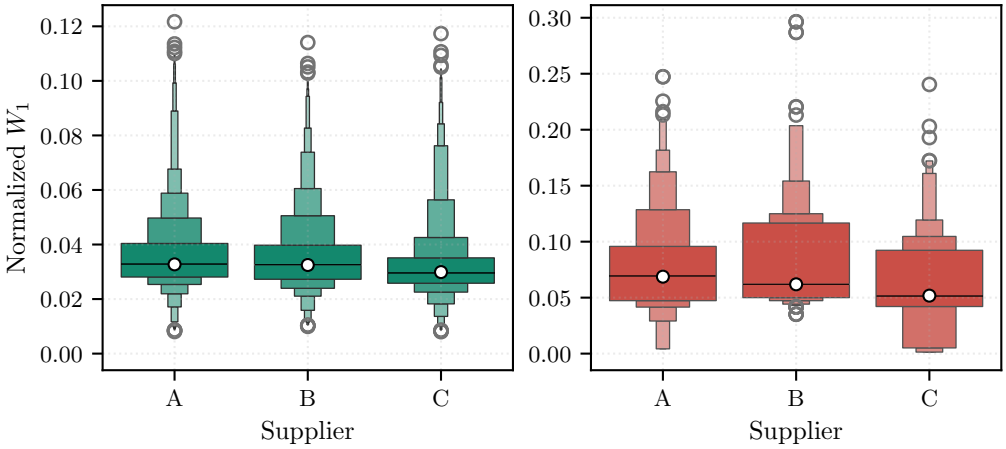

Figure 5.2: Range-normalized approximation error of the LCA coefficient reduction, shown as $W_1/(\max - \min)$ across all non-constant coefficients. Results are grouped by supplier and by impact dimension (environmental vs. economic). Each observation corresponds to one process coefficient under one scenario-tree realization (seed).

Table 5.1 provides a portfolio-level summary of the same normalized error metric (in % of each coefficient's outcome range). Across product categories, environmental approximation errors are relatively homogeneous, whereas economic approximation errors vary more strongly and are driven by a small number of difficult-to-approximate categories/coefficients. Differences between suppliers are portfolio-conditional: each supplier is associated

with a different set of underlying coefficient distributions, yet no supplier shows consistently elevated errors across the portfolio. The dominant contrast is therefore between environmental and economic dimensions rather than between suppliers.

Table 5.1: Range-normalized W_1 distances ($100 \cdot W_1/(\max - \min)$, reported in % of each coefficient's outcome range) by supplier and product category (A–C). Table entries are medians over all coefficient–seed samples within each (product category, supplier) subset; 'Total' pools suppliers within a product category. 'Median (all categories)' is the median of the product-category medians (equal weight per product category). 'Median (all trees)' is the median over all coefficient–seed samples across the full portfolio, reported per supplier and pooled ('Total'). Some product categories do not exhibit cost variation, for which (–) is reported.

	Environmental [%]				Economic [%]			
Product category	A	B	C	Total	A	B	C	Total
Conference chairs	4.63	3.40	2.50	3.54	8.60	10.76	6.72	8.23
Conference tables	3.05	2.95	2.96	2.98	4.73	4.73	4.73	4.73
Education chairs	2.97	3.95	2.51	3.09	0.43	–	0.14	0.28
Education tables	3.19	2.84	2.90	2.97	5.00	5.00	5.00	5.00
Electric office desks	3.67	3.43	3.29	3.47	5.01	5.27	6.22	5.62
Manual office desks	3.14	3.30	3.12	3.18	4.61	5.50	4.21	5.50
Office chairs	3.67	4.23	3.39	3.84	8.05	7.57	7.12	7.68
Pedestals	2.92	2.96	2.99	2.95	–	–	–	–
Storage cabinets	2.60	2.60	2.74	2.66	–	–	–	–
Median (all categories)	3.14	3.30	2.96	3.09	5.00	5.38	5.00	5.50
Median (all trees)	3.27	3.25	2.99	3.16	6.87	6.19	5.18	6.09

5.1.2 Demand and Returns

Figure 5.3 compares empirical and tree-implied moments for all (stage, category) pairs. Empirical means align closely with the median tree-implied means (left panel), indicating that the tree construction preserves the mean structure with little bias. In contrast, tree-implied standard deviations are systematically smaller than empirical benchmarks (right panel), implying structural under-dispersion, especially for pairs with large empirical variability. Consequently, the scenario approximation represents typical operating conditions more accurately than extreme tail events, for the risk-neutral optimization considered here, the value of hedging against rare high-demand peaks may be under-represented. Moment-based metrics are preferred over aggregated distributional distance metrics (e.g., Wasserstein distance) because the resulting mismatch patterns provide direct, decision-relevant diagnostics.

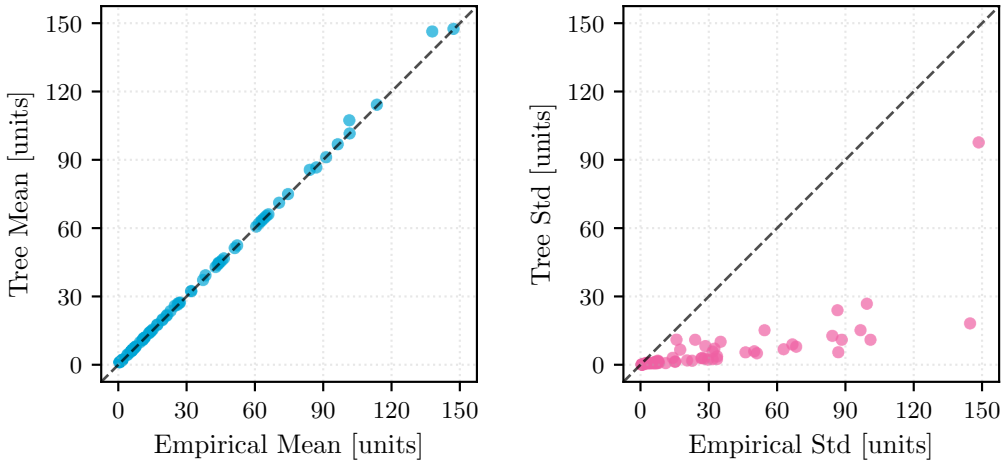

Figure 5.3: Empirical versus tree-implied demand moments across (stage, product-category) pairs (Stage 1 excluded). Left: empirical mean versus the median tree mean across 30 scenario trees. Right: empirical standard deviation versus the median tree standard deviation across 30 scenario trees. The dashed line indicates perfect agreement ($y = x$).

To quantify these deviations, Figure 5.4 reports stage-wise normalized root mean squared error (NRMSE) distributions across the 30 trees. Mean errors are consistently small across stages. By contrast, NRMSE values for standard deviations are large at every stage (often close to unity in relative terms), consistent with the moment-scatter diagnostic. Tree-to-tree variation is present but secondary compared to the overall dispersion shortfall.

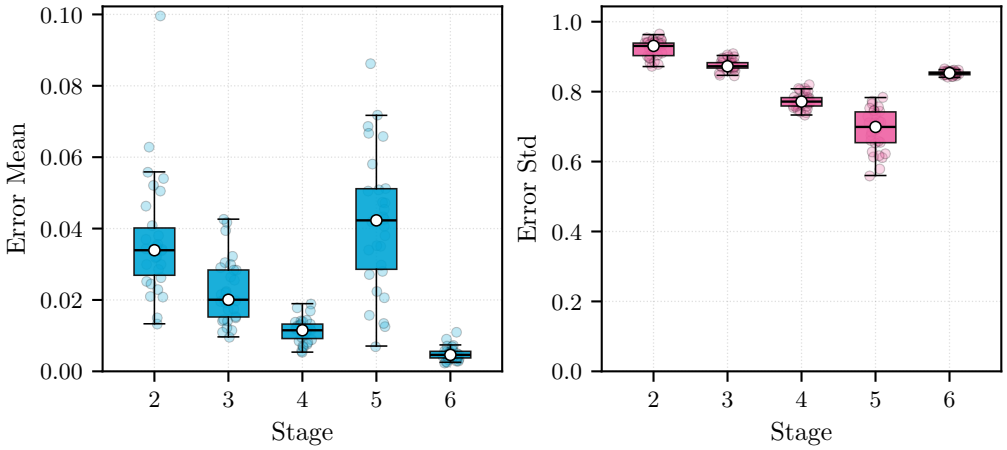

Figure 5.4: Distribution of stage-wise relative RMS error (NRMSE) across 30 scenario trees (Stage 1 excluded). For each stage, the error is computed as RMSE across product categories between empirical and tree-implied moments, normalized by the empirical RMS across product categories at that stage. White markers show the median across trees; vertical lines show the min–max across trees. Left: errors for the mean vector. Right: errors for the standard deviation vector. Note the y-axis scale difference.

Table 5.2: Category-level NRMSE of the mean demand by stage (%; Stage 1 excluded). Values in the table body use the average tree (mean across 30 scenario trees). The 'Total' column is the RMS aggregation across stages normalized by empirical RMS. The 'Median (tree dist.)' row reports the median system-level NRMSE of individual trees.

Product category	Stage 2	Stage 3	Stage 4	Stage 5	Stage 6	Total
Conference chairs	1.14	0.30	0.42	1.25	0.28	0.76
Conference tables	0.22	1.62	0.35	0.46	0.18	1.16
Education chairs	0.39	0.48	0.02	6.15	0.16	4.77
Education tables	0.45	0.32	0.32	4.93	0.12	4.63
Electric office desks	0.43	0.19	1.02	0.05	0.23	0.64
Manual office desks	0.71	0.14	0.19	0.28	0.38	0.30
Office chairs	0.04	0.27	0.01	0.10	0.43	0.21
Pedestals	0.25	1.56	2.01	0.26	0.10	1.33
Storage cabinets	0.21	2.78	1.25	0.25	0.06	1.82
Conference chairs (R)	0.45	0.75	0.06	0.15	0.05	0.38
Conference tables (R)	23.74	28.28	25.00	22.09	19.92	23.85
Education chairs (R)	0.01	0.02	0.04	0.00	0.01	0.02
Education tables (R)	0.34	0.18	0.13	0.07	0.04	0.18
Electric office desks (R)	0.07	0.43	0.45	0.74	0.23	0.45
Manual office desks (R)	0.56	0.18	0.00	0.03	0.04	0.25
Office chairs (R)	0.03	0.03	0.02	0.00	0.03	0.02
Pedestals (R)	0.45	0.13	0.62	0.33	0.05	0.37
Storage cabinets (R)	6.66	1.83	1.58	1.48	1.95	3.24
Median (all categories)	0.41	0.31	0.34	0.27	0.14	0.54
Median (tree dist.)	3.39	2.01	1.15	4.23	0.46	3.02

Table 5.3: Category-level NRMSE of the standard deviation of demand by stage (%; Stage 1 excluded). Values compare empirical standard deviations to the mean of the tree-implied standard deviations across 30 scenario trees. High NRMSE values indicate substantial mismatch in dispersion (tree variance compression).

Product category	Stage 2	Stage 3	Stage 4	Stage 5	Stage 6	All
Conference chairs	92.37	88.93	84.49	87.11	88.41	86.71
Conference tables	92.53	31.48	88.51	89.84	86.80	59.21
Education chairs	93.07	89.25	92.80	68.03	90.99	70.34
Education tables	94.22	91.76	91.89	33.99	90.56	39.80
Electric office desks	92.76	85.07	62.95	91.12	55.19	62.88
Manual office desks	91.68	90.19	72.96	88.71	86.74	83.15
Office chairs	92.27	89.49	72.69	87.40	84.79	82.59
Pedestals	87.42	78.98	72.28	89.20	82.16	79.32
Storage cabinets	89.12	71.92	72.43	89.58	79.34	77.42
Conference chairs (R)	79.58	84.84	79.16	81.59	79.15	80.86
Conference tables (R)	89.47	100.00	99.35	99.08	92.80	96.26
Education chairs (R)	77.43	83.96	84.16	86.08	82.73	82.95
Education tables (R)	79.41	85.40	80.60	83.74	80.81	82.02
Electric office desks (R)	76.17	85.75	80.59	84.03	79.38	81.27
Manual office desks (R)	81.11	84.35	81.53	83.55	80.16	82.14
Office chairs (R)	81.89	84.37	83.17	86.13	82.25	83.59
Pedestals (R)	76.43	81.23	79.28	79.57	78.20	78.96
Storage cabinets (R)	89.08	96.85	88.66	93.75	81.34	89.98
Median (all categories)	89.10	85.23	81.07	86.62	82.21	81.65
Median (tree dist.)	93.07	87.26	77.15	69.90	85.31	73.42

Tables 5.2 and 5.3 disaggregate errors by product category. For forward demand, mean-demand errors are typically well below a few percent across stages, with larger deviations concentrated in a small number of category–stage combinations (notably Stage 5 for the education category). Return means are also generally reproduced well, validating the cohort-based reliability model (Section 4.5.5), with some exceptions such as conference tables (R), where relative errors are large across all stages. The standard-deviation table confirms that for both forward and return flows, NRMSE values are high across most categories and stages, suggesting that the second-moment mismatch is structural and the scenario approximation method fails to conserve variability.

5.2 Optimization Outcomes

Figure 5.5 summarizes the resulting portfolio of solutions in the environmental–economic plane. Each marker corresponds to one (λ, supplier) combination averaged over seeds, with colors and shapes indicating the supplier. The horizontal axis reports the expected life cycle GHG emissions in tCO_2eq, and the vertical axis reports the expected total cost in k€. Solutions that are not dominated in either dimension form an empirical Pareto front, highlighted by the dashed line.

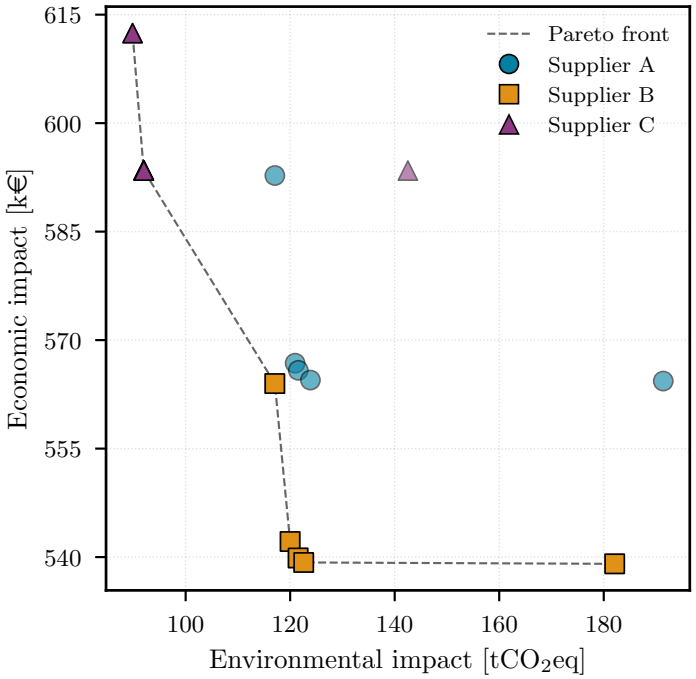

Figure 5.5: Environmental–economic trade-off of optimized supplier portfolios. Each point represents one optimized solution for a given weight λ and supplier (A–C), aggregated over 30 scenario-tree seeds. The horizontal axis shows expected environmental impact (tCO_2eq); the vertical axis shows expected economic impact (k€). Colors and marker shapes distinguish suppliers; the dashed line connects non-dominated solutions and approximates the Pareto front.

The Pareto front in Figure 5.5 is relatively short and steep, indicating an asymmetry between the two objectives. Moving from the most environmentally favorable configuration (low emissions, higher cost) towards the most economical configuration (lowest cost, higher emissions) reduces costs by only a relatively modest amount, while almost doubling GHG emissions. The efficient frontier is spanned primarily by solutions associated with suppliers B and C, whereas solutions associated with supplier A are dominated once both objectives are considered jointly.

Table 5.4 provides a consolidated numerical summary of the environmental and economic outcomes for each supplier across the objective weights λ, reporting mean performance over the 30 scenario-tree seeds. Consistent with Figure 5.5, supplier C achieves the lowest GHG emissions for all λ, whereas supplier B achieves the lowest total cost for all λ. Under full environmental minimization ($\lambda = 0$), supplier C yields approximately 90 tCO_2eq at 612 k€, while supplier B yields 117 tCO_2eq at 564 k€. Under full economic minimization ($\lambda = 1$), supplier B attains the lowest cost of about 539 k€, but with emissions increasing to roughly 182 tCO_2eq, while supplier C remains the lowest-emission option at 143 tCO_2eq at 593 k€. Across the range $\lambda \in [0, 0.75]$, costs decrease by several tens of k€ while emissions increase only marginally (a few tCO_2eq); the largest increase in emissions occurs at $\lambda = 1.00$, with comparatively limited additional cost reduction relative to $\lambda = 0.75$.

Table 5.4: Mean economic and environmental performance over 30 seeds by supplier under varying objective weight λ. Here, $\lambda = 0$ corresponds to full environmental minimization, while $\lambda = 1$ corresponds to full economic minimization. Values are rounded to the nearest integer. Bold entries indicate the minimum (best) value within each row for each block (Environmental and Economic).

	Environmental (tCO_2eq)			Economic (k€)		
λ	A	B	C	A	B	C
0.00	117	117	**90**	593	**564**	612
0.25	121	120	**92**	567	**542**	594
0.50	122	122	**92**	566	**540**	593
0.75	124	123	**92**	564	**539**	593
1.00	191	182	**143**	564	**539**	593

To characterize how the optimized portfolios change internally as λ increases, Table 5.5 reports signed normalized step changes by technology category across consecutive λ intervals. Each entry is normalized by the supplier-specific median level across runs and all λ values, so these values should be interpreted as direction and relative magnitude of reallocation between steps (not as a percent change from a baseline).

Table 5.5: Signed normalized step changes. Each entry is $100 \cdot (\Delta x / \tilde{x})$, where Δx is the net change over the indicated λ step and $\tilde{x}$ is the supplier-specific median level computed across runs and all λ. Positive values indicate net shifts toward the technology; negative values indicate shifts away. For absolute values, see Table B.3.

		λ step			
Supplier	Technology	0.00→0.25	0.25→0.50	0.50→0.75	0.75→1.00
A	New production	14.2%	1.1%	2.6%	**83.9%**
A	Refurbishment	−15.7%	−1.0%	−2.5%	**−78.8%**
A	Revitalization	4.1%	0.0%	0.0%	−3.4%
A	Warehousing	**−94.4%**	**−54.9%**	**−81.2%**	−8.0%
A	Diesel transport	0.0%	0.0%	0.0%	895.4%[b]
A	Electric transport	0.0%	0.0%	0.0%	**−97.0%**
A	Disposal	−4.2%	0.0%	0.0%	3.5%
B	New production	13.3%	6.7%	2.1%	**78.6%**
B	Refurbishment	−15.4%	−6.9%	−2.1%	**−79.0%**
B	Revitalization	3.3%	0.2%	−0.1%	−2.4%
B	Warehousing	**−71.7%**	**−53.8%**	**−63.1%**	−27.6%
B	Diesel transport	0.0%	0.0%	0.0%	3,859.2%[b]
B	Electric transport	0.0%	0.0%	0.0%	**−100.2%**
B	Disposal	−3.6%	−0.2%	0.1%	2.6%
C	New production	17.7%	0.2%	0.0%	**86.9%**
C	Refurbishment	−19.5%	−0.2%	0.0%	**−84.2%**
C	Revitalization	3.9%	0.0%	0.0%	−2.8%
C	Warehousing[a]	**−38.6%**	−0.7%	−0.1%	751.2%[c]
C	Diesel transport	0.0%	0.0%	0.0%	3,859.2%[b]
C	Electric transport	0.0%	0.0%	0.0%	**−100.2%**
C	Disposal	−4.2%	0.0%	0.0%	3.0%

[a] Warehousing values affected by a run with zero cost.
[b] Very large values occur when the typical level $\tilde{x}$ for that category is near zero (small denominator), so the ratio can exceed 100% substantially.
[c] Same caveat as (b), plus the data-quality issue noted in (a).

Across suppliers, the step-change results show a consistent reallocation away from

refurbishment and toward new production as the objective shifts toward cost minimization. The direction of this shift is already visible in the first interval ($0.00 \rightarrow 0.25$), with signed normalized increases in new production (approximately 13%–18%) and corresponding decreases in refurbishment (approximately -15%–-20%). The most pronounced stepwise reallocation occurs in the final interval ($0.75 \rightarrow 1.00$), where new production increases sharply (79%–87%) and refurbishment decreases comparably (-79%–-84%) for suppliers A–C. Over the intermediate intervals ($0.25 \rightarrow 0.75$), most categories exhibit comparatively small changes, indicating broadly stable decision mixes between consecutive weights until the final shift to $\lambda = 1.00$. Warehousing exhibits large negative step changes for suppliers A and B across multiple intervals, while supplier C displays an anomalous final-step increase, consistent with the data-quality caveats noted in Table 5.5. Transport-category ratios are concentrated in the final interval: electric transport decreases strongly, while diesel transport shows very large positive ratios; as noted in the table footnotes, these extreme ratios primarily reflect near-zero typical levels (small denominators) rather than large absolute changes.

Finally, Table 5.6 reports the fraction of exogenous input requirements that exceed the capacity constraint thresholds, by product category and supplier. These exceedance rates are invariant across λ; they are driven by the exogenous demand structure rather than the objective weighting used in the portfolio optimization.

Table 5.6: Percentage of exogenous inputs exceeding the constraint threshold by supplier, grouped by constraint type. Values are in %, rounded to two decimals. These values apply to all λ values.

	> Circularity [%]			> Electric delivery [%]		
Product category	A	B	C	A	B	C
Conference chair	3.23	0.53	0.18	10.33	0.03	0.03
Conference table	5.58	5.57	5.51	5.58	4.21	4.21
Education chair	–	–	–	7.02	3.06	3.06
Education table	6.74	6.50	6.42	7.47	6.25	6.25
Electric office desk	2.86	1.38	1.02	5.79	0.53	0.53
Manual office desk	9.54	6.26	5.07	18.38	2.46	2.46
Office chair	7.88	4.80	3.53	14.32	0.59	0.59
Pedestal	–	–	–	3.98	0.11	0.11
Storage cabinet	–	–	–	29.19	0.53	0.53

Exceedance rates for the circularity threshold (where defined) are generally in the single digits, with the highest values observed for manual office desks (up to 9.54% for supplier A) and office chairs (up to 7.88% for supplier A), while conference and education tables exhibit consistent exceedance rates around 5.5%–6.7% across suppliers. For the electric-delivery threshold, supplier A exhibits substantially higher exceedance rates across multiple categories (e.g., storage cabinets 29.19%, manual office desks 18.38%, office chairs 14.32%, conference chairs 10.33%), whereas suppliers B and C are near zero for several categories (often below 1%), with the main exceptions occurring for table and education categories. Operationally, these exceedances identify product categories and supplier configurations for which exogenous requirements cannot be met within the constrained channel, implying the need for compensating measures such as additional sourcing of new items and/or buffering via inventory/storage.

Overall, the optimization outcomes indicate that supplier C dominates on expected emissions and supplier B dominates on expected cost across the evaluated weights, while the internal technology mix shifts most strongly between $\lambda = 0.75$ and $\lambda = 1.00$. The next section assesses whether these differences are statistically distinguishable from the SAA benchmark using paired optimality-gap estimates.

5.3 Optimality Gap

Figure 5.6 presents the paired relative optimality gaps ($\hat{G}_{\text{rel}}$). The analysis evaluates, for a given objective weight λ and a supplier decision $\hat{x}$, the magnitude of the expected performance loss relative to the best achievable SAA objective and determines whether this loss is statistically distinguishable from zero under the adopted replication design.

The paired Mak–Morton–Wood estimator compares, for each replication, the objective value of a fixed supplier decision $\hat{x}$ upper bound (UB) against the objective value of the unrestricted SAA solution computed on the same evaluation scenario tree lower bound (LB). Using the same tree realization for UB and LB yields paired differences with reduced variance. Relative gaps are computed as $\hat{G}_{\text{rel}} = \hat{G}/\hat{\mu}_{\text{LB}}$ and reported in percent; at $\lambda = 0$ and $\lambda = 1$ this corresponds to relative performance loss in emissions and cost units, respectively, while for $0 < \lambda < 1$ it measures loss in the weighted index.

Figure 5.6 visualizes the estimated relative gaps $\hat{G}_{\text{rel}}$ across weights λ and suppliers, including replication-level variability (faint markers) and 95% confidence intervals (error bars). Table 5.7 reports the corresponding numerical estimates: the sample means of UB and LB objective values, the paired gap estimate $\hat{G}$, and 95% t-intervals, both in absolute and relative terms.

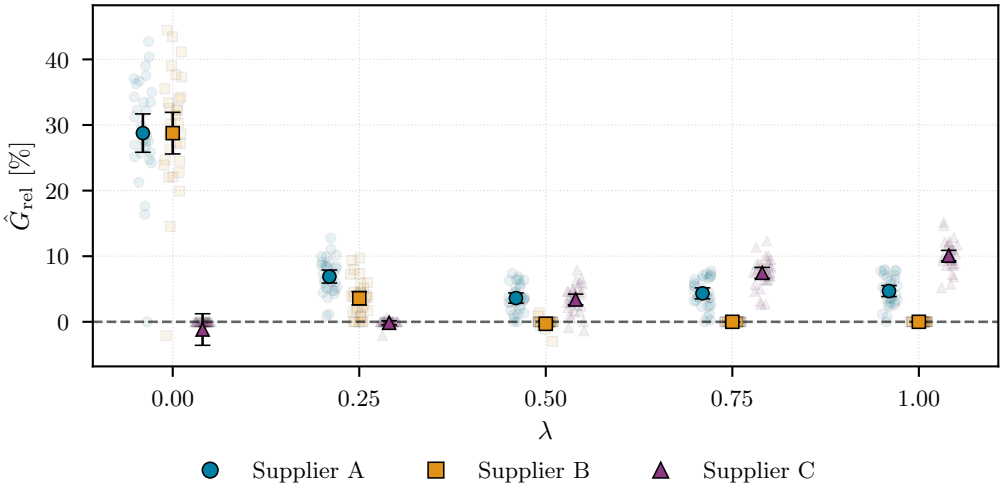

Figure 5.6: Mak–Morton–Wood paired relative optimality-gap estimates $\hat{G}_{\text{rel}}$ across objective weights λ and suppliers. Faint markers show replication-level paired gaps (in percent); error bars show the mean estimate with a 95% t-interval over $n = 30$ paired scenario-tree replications. The horizontal dashed line indicates a zero gap; values near zero indicate that the supplier is statistically indistinguishable from the best achievable SAA objective under the same scenario approximation.

Table 5.7: Mak–Morton–Wood paired gap estimates for each locked supplier $\hat{x}$ and weight λ over $n = 30$ paired scenario-tree replications. Objective values refer to the weighted objective $h(\cdot;\lambda)$. At $\lambda = 0$, h equals environmental impact [tCO_2eq]; at $\lambda = 1$, h equals economic impact [k€]; for $0 < \lambda < 1$ it is a dimensionless weighted index. Confidence intervals are 95% t-intervals. Small negative gaps may occur due to mixed-integer solver tolerances.

		Objective value					Relative gap [%]		
Supplier	λ	$\hat{\mu}_{UB}$	$\hat{\mu}_{LB}$	$\hat{G}$	CI_{lo}	CI_{hi}	$\hat{G}_{rel}$	CI_{lo}	CI_{hi}
A	0.00	117,071	90,916	26,156	23,495	28,816	28.77	25.84	31.70
	0.25	232,415	217,416	14,999	12,851	17,146	6.90	5.91	7.89
	0.50	343,691	331,694	11,997	9,423	14,571	3.62	2.84	4.39
	0.75	454,326	435,437	18,888	15,268	22,509	4.34	3.51	5.17
	1.00	564,338	539,065	25,273	20,755	29,790	4.69	3.85	5.53
B	0.00	117,065	90,916	26,149	23,266	29,032	28.76	25.59	31.93
	0.25	225,529	217,706	7,823	5,593	10,053	3.59	2.57	4.62
	0.50	330,700	331,716	−1,016	−2,358	326	−0.31	−0.71	0.10
	0.75	435,092	435,092	0	0	0	0.00	0.00	0.00
	1.00	539,068	539,065	3	−3	8	0.00	0.00	0.00
C	0.00	89,846	90,916	−1,069	−3,257	1,118	−1.18	−3.58	1.23
	0.25	217,357	217,691	−334	−988	319	−0.15	−0.45	0.15
	0.50	342,720	331,446	11,275	8,558	13,991	3.40	2.58	4.22
	0.75	468,082	435,602	32,480	28,825	36,135	7.46	6.62	8.30
	1.00	593,444	539,065	54,378	50,074	58,683	10.09	9.29	10.89

Three patterns are observed. First, at the environmental extreme ($\lambda = 0$), supplier C is statistically indistinguishable from the best achievable SAA objective. In contrast, locking to suppliers A or B induces a large and clearly positive optimality gap of approximately 29 %, corresponding to an absolute loss of about 26 tCO_2eq. This indicates that, under pure environmental minimization, supplier selection is a main driver of performance.

Second, at the economic extreme ($\lambda = 1$), supplier B is effectively optimal within sampling and solver tolerances (a near-zero estimated gap), whereas supplier C becomes suboptimal with an estimated relative gap of about 10 %, corresponding to an economic penalty of approximately 54 k€. Supplier A remains significantly worse than the benchmark, with a relative gap of roughly 5 %.

Third, in the intermediate regime ($\lambda \in \{0.25, 0.50, 0.75\}$), supplier B is near-optimal for $\lambda \geq 0.50$ (gaps at or near zero), while supplier C transitions from near-optimal at $\lambda \leq 0.25$ to clearly suboptimal as λ increases. The zero gaps for supplier B at $\lambda \geq 0.75$ indicate that the locked supplier decision coincides with the unrestricted SAA optimum in all paired replications (up to solver tolerance). Supplier A exhibits consistently positive gaps across all intermediate weights, indicating that it is rarely selected by the unrestricted SAA benchmark.

In summary, the SAA benchmark selects supplier C at low λ and switches to supplier B as λ increases, with supplier A remaining uncompetitive across the entire range. This confirms that the primary strategic choice lies between the environmentally superior supplier C and the cost-effective supplier B.

5.4 Value of Information

Estimates for the EVPI are positive for most λ values (Figure 5.7). This section quantifies two related aspects of the case study introduced in Chapter 4. First, it quantifies the potential performance improvement that would be achieved if all uncertain quantities (demand/returns and supplier-dependent impact coefficients) were revealed before any decisions are taken, allowing scenario-contingent re-optimization of the full decision set

including supplier selection, as measured by the EVPI. Second, it assesses whether explicitly solving the multistage stochastic model yields a measurable improvement relative to an expected-value approximation with recourse, as measured by the VSS.

All estimates are based on $R = 30$ independent scenario-tree replications. In each replication, RP, WS, and expected value of the expected value (EEV) decisions are computed and evaluated on the same scenario-tree realization to preserve pairing and reduce estimator variance. For each λ, all models use the same weighted objective $h(\cdot;\lambda)$ and differ only in their information structure. Results are reported as outcome-based measures in terms of total emissions and total cost under the corresponding decisions. WS assumes full revelation of all uncertainties at the initial stage (perfect foresight over the entire horizon).

Figure 5.7 visualizes replication-wise EVPI values as a function of the trade-off parameter λ, together with the corresponding mean and 95% t-intervals. Intervals assume replication-wise independence across scenario-tree seeds. Table 5.8 reports the underlying outcome values ($\hat{z}_{\mathrm{rp}}$, $\hat{z}_{\mathrm{ws}}$, $\hat{z}_{\mathrm{eev}}$) and summarizes the estimated EVPI and VSS statistics. Note that EVPI and VSS are computed from unrounded outcome estimates; rounding applies only to displayed outcome values.

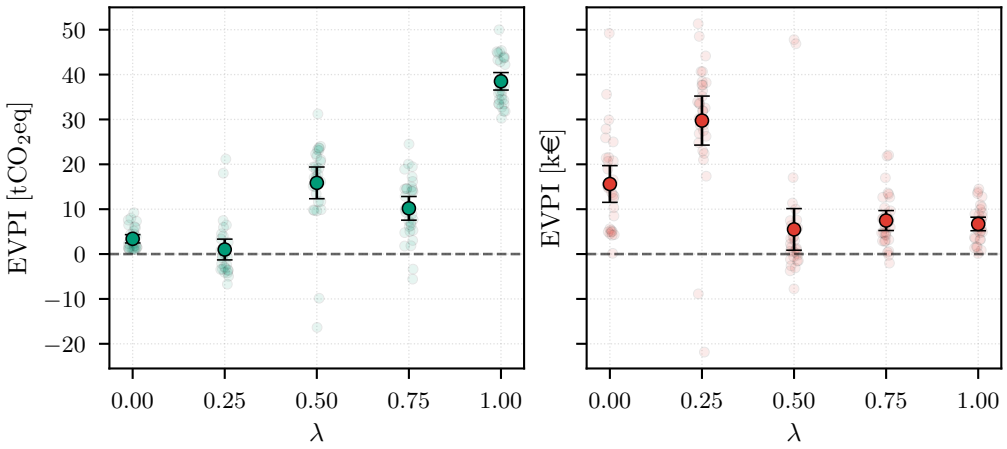

Figure 5.7: EVPI estimated from $R = 30$ independent replications, shown as a function of the trade-off parameter λ. Faint points show replication-wise estimates; markers with error bars show the sample mean and 95% t-interval. Left: EVPI expressed in environmental units (tCO_2eq). Right: EVPI expressed in economic units (k€). Small negative realizations may occur due to finite-sample variability and mixed-integer solver tolerances and are interpreted as noise around zero.

The corresponding VSS estimates are shown in Figure 5.8. In relative terms, the performance gain from the stochastic solution is negligible (relative VSS $< 1\%$) across all λ, indicating that for the evaluated furniture portfolio the strategic supplier selection is robust to the modeled uncertainty: the supplier that is optimal under expected conditions remains optimal when stochasticity is explicitly represented.

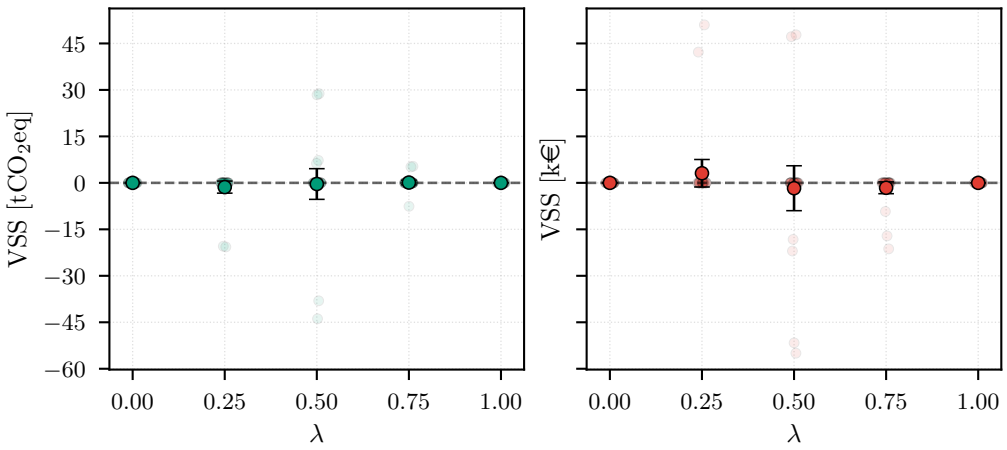

Figure 5.8: Value of the Stochastic Solution (VSS) estimated from $R = 30$ independent replications. Faint points show replication-wise estimates; markers with error bars show the sample mean and 95% t-interval. Left: VSS in environmental units (tCO_2eq). Right: VSS in economic units (k€). Values near zero indicate that the expected-value approximation performs comparably to the fully stochastic formulation in expected terms for the evaluated instance.

Table 5.8: Value of information analysis over $R = 30$ independent replications. Outcome values ($\hat{z}_{rp}$, $\hat{z}_{ws}$, $\hat{z}_{eev}$) are scaled by 10^3 (Environmental: tCO_2eq; Economic: k€) and rounded to whole numbers. EVPI denotes the expected value of perfect information and VSS the value of the stochastic solution. Confidence intervals are 95% t-intervals.

	Outcome values			EVPI					VSS				
λ	$\hat{z}_{rp}$	$\hat{z}_{ws}$	$\hat{z}_{eev}$	$\hat{\mu}$	$\hat{\sigma}$	SE	CI_{lo}	CI_{hi}	$\hat{\mu}$	$\hat{\sigma}$	SE	CI_{lo}	CI_{hi}
Environmental impact [tCO_2eq]													
0.00	90	86	90	3.40	2.47	0.45	2.48	4.32	0.00	0.00	0.00	0.00	0.00
0.25	93	92	92	1.01	6.20	1.13	−1.31	3.32	−1.37	5.22	0.95	−3.32	0.58
0.50	120	105	120	15.86	9.47	1.73	12.33	19.40	−0.37	13.26	2.42	−5.32	4.58
0.75	124	113	124	10.19	7.06	1.29	7.55	12.82	0.10	1.97	0.36	−0.64	0.83
1.00	183	144	183	38.50	5.22	0.95	36.55	40.45	0.00	0.00	0.00	0.00	0.00
Economic impact [k€]													
0.00	612	596	612	15.62	10.96	2.00	11.52	19.71	0.00	0.00	0.00	0.00	0.00
0.25	590	560	593	29.75	14.63	2.67	24.28	35.21	3.11	11.88	2.17	−1.33	7.54
0.50	545	539	543	5.50	12.41	2.27	0.86	10.13	−1.73	19.41	3.54	−8.98	5.52
0.75	541	534	539	7.47	5.95	1.09	5.25	9.69	−1.59	5.11	0.93	−3.50	0.32
1.00	539	532	539	6.71	4.04	0.74	5.20	8.21	0.00	0.00	0.00	0.00	0.00

Across both outcome dimensions, EVPI is positive for most λ, and for several weights the 95% intervals lie strictly above zero. For example, at $\lambda = 1$ the emissions EVPI is 38.5 tCO_2eq, corresponding to roughly 21% of the emissions outcome under the RP decision. At $\lambda = 0.25$, the cost EVPI is ≈ 29.8 k€ (about 5% of $\hat{z}_{rp}$). In contrast, VSS is small and statistically indistinguishable from zero across λ.

5.5 Uncertainty of Outcomes

5.5.1 Contribution analysis

This section characterizes the distribution of realized environmental and economic impacts under uncertainty using an ex-post policy evaluation design. For each objective weight λ

and scenario-tree replication (seed), the scenario-tree model yields a scenario-contingent scaling policy $S_{\lambda,r,s}$ (process amounts) and an associated supplier choice. In the ex-post evaluation, these decisions are held fixed. Each evaluation draw then (i) samples a scenario node s from the scenario tree according to its probability (and a seed r uniformly over replications), and (ii) samples an evaluation draw of the uncertain process impact coefficients $\tilde{a}$. Realized impacts are computed by multiplying the fixed process amounts with the sampled coefficients and summing across processes, i.e., $Y_\lambda = \sum_p S_{\lambda,r,s}(p)\,\tilde{a}(p)$, and are subsequently aggregated to supplier totals and decomposed into stage-wise and product-category contributions. Unless stated otherwise, results are pooled over objective weights $\lambda \in \{0, 0.25, 0.5, 0.75, 1\}$, yielding a mixture over preference weights and therefore combining uncertainty-driven outcome variability with systematic differences in policies induced by λ. The analysis focuses on (i) the distribution of total impacts by supplier and (ii) the decomposition of these totals into stage-wise and product-category contributions.

Figure 5.9 reports the distribution of total impacts for each supplier, shown separately for environmental and economic outcomes. Table 5.9 summarizes the corresponding descriptive statistics. For environmental impacts, Suppliers A and B exhibit nearly identical central tendency (means 138.56 and 138.33 tCO_2eq), whereas Supplier C shows a substantially lower level (mean 107.18 tCO_2eq) and a lower upper tail. For economic impacts, Supplier B achieves the lowest typical cost (mean 563.65 k€), while Supplier C shows the highest typical cost (mean 611.86 k€) and the largest dispersion (standard deviation 68.47 k€), with a larger right tail (maximum 1,085.8 k€). These results indicate that Supplier C is associated with lower environmental burden but higher and more variable economic outcomes, whereas Supplier B is economically favorable but environmentally comparable to Supplier A.

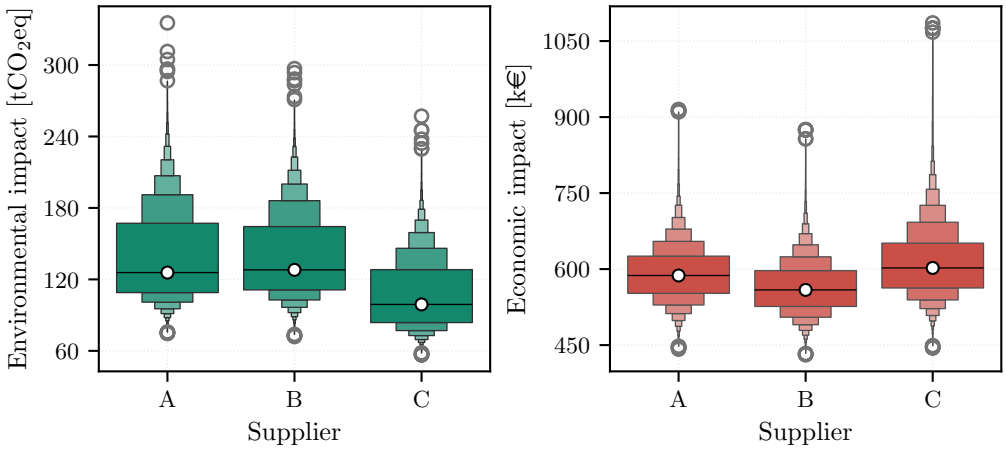

Figure 5.9: Distributions of total impacts by supplier (pooled over λ). Left: environmental impact [tCO_2eq]. Right: economic impact [k€]. Boxenplots show progressively deeper quantiles; the white dot marks the median.

The pooled distributions mask substantial policy-induced shifts in the level of impacts, especially for emissions at $\lambda = 1$. Detailed summary statistics disaggregated by λ are provided in Appendix B (Table B.1). In contrast, Table 5.9 reports the main aggregated results.

Table 5.9: Summary statistics for environmental and economic impacts by supplier (n = 50,000 samples per supplier).

	Descriptive statistics						
Supplier	μ	σ	min	$Q_{0.25}$	$Q_{0.50}$	$Q_{0.75}$	max
Environmental impact [tCO_2eq]							
A	138.56	38.29	74.69	108.91	125.73	167.13	335.34
B	138.33	35.06	72.02	111.19	128.00	164.37	296.97
C	107.18	29.40	56.73	83.74	98.99	128.21	257.05
Economic impact [k€]							
A	591.12	55.31	442.03	552.19	587.38	625.62	914.80
B	563.65	52.20	430.94	526.28	558.90	597.00	875.32
C	611.86	68.47	444.82	562.75	602.27	651.22	1,085.80

Figure 5.10 and Table 5.10 decompose total impacts into four life cycle stages by reporting relative contributions (shares of the total impact per sample). Environmental impacts are overwhelmingly driven by manufacturing (mean share 96.51%), while transportation and end-of-life contribute on average 1.03% and 2.44%, respectively; warehousing is negligible (mean 0.02%). Despite the high concentration in manufacturing, the end-of-life share shows the heaviest upper tail (maximum 13.18%), indicating that in a minority of realizations end-of-life processes account for a substantially larger fraction of total environmental impact. Economic impacts are also dominated by manufacturing (mean share 87.54%), but transportation (8.93%) and warehousing (3.54%) represent materially larger fractions than in the environmental case. The end-of-life share for economic impact is identically zero under the present cost accounting.

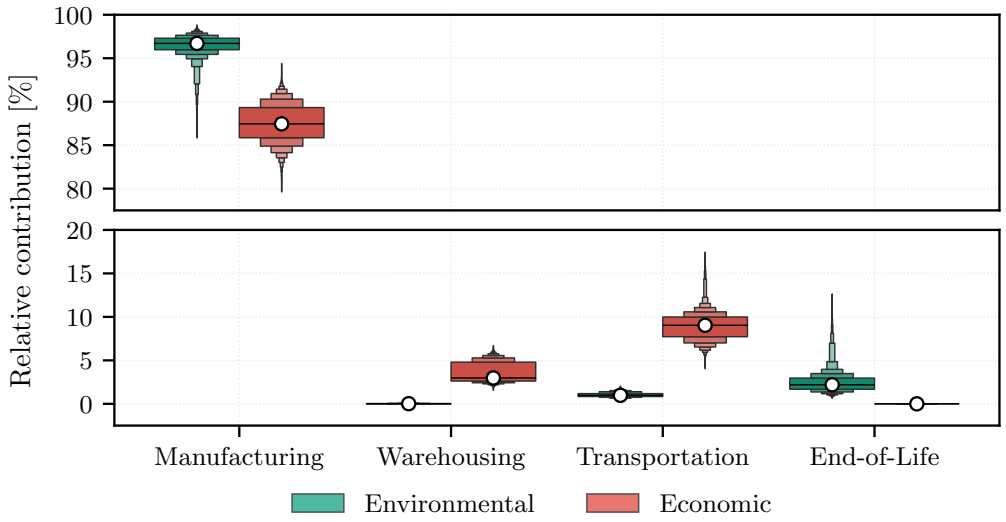

Figure 5.10: Stage-wise relative contribution shares (pooled over suppliers and λ). Boxenplots summarize the distribution of contribution shares for environmental and economic impacts. Each observation corresponds to one evaluation draw.

Table 5.10: Summary statistics for stage-wise contributions (in %) grouped by impact type ($n = 150{,}000$ samples per stage).

	Descriptive statistics						
Stage	μ	σ	min	$Q_{0.25}$	$Q_{0.50}$	$Q_{0.75}$	max
Environmental impact share [%]							
Manufacturing	96.51	1.24	85.35	95.98	96.71	97.30	98.88
Warehousing	0.02	0.02	0.00	0.00	0.02	0.03	0.29
Transportation	1.03	0.28	0.40	0.84	0.98	1.18	2.06
End-of-Life	2.44	1.21	0.59	1.68	2.19	2.96	13.18
Economic impact share [%]							
Manufacturing	87.54	2.24	79.33	85.85	87.45	89.33	94.49
Warehousing	3.54	1.19	1.52	2.62	2.98	4.80	6.84
Transportation	8.93	1.60	3.98	7.72	9.03	9.99	17.67
End-of-Life	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Figure 5.11 and Table 5.11 report relative contributions by product category. Environmental impacts are primarily driven by Storage cabinets (mean share 42.72%), followed by Office chairs (12.94%) and Pedestals (12.83%). Several categories exhibit substantial dispersion, most notably Office chairs (standard deviation 11.32%) and Storage cabinets (standard deviation 10.05%). Economic impacts are more evenly distributed across categories, with Manual office desks (18.51%), Office chairs (18.31%), and Storage cabinets (16.05%) accounting for the largest mean shares. Compared to the environmental case, the leading economic categories have tighter distributions (e.g., Manual office desks: standard deviation 1.95%).

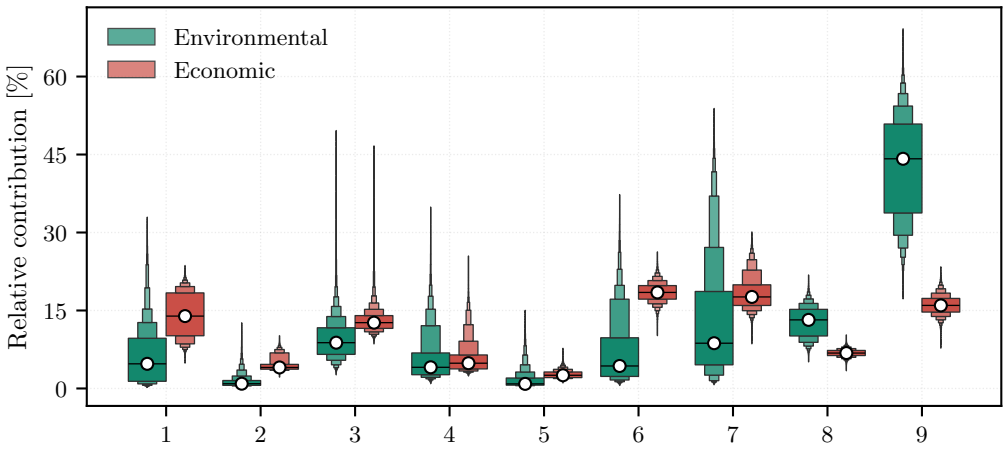

Figure 5.11: Product-category relative contribution shares (pooled over suppliers and λ). Distributions are shown separately for environmental and economic impacts. Each observation corresponds to one evaluation draw. Numeric labels denote product categories: (1) Conference chairs, (2) Conference tables, (3) Education chairs, (4) Education tables, (5) Electric office desks, (6) Manual office desks, (7) Office chairs, (8) Pedestals, and (9) Storage cabinets.

To complement the outcome distributions, Table 5.12 reports the coefficient of variation of the replication-level mean outcome estimates across scenario-tree seeds ($R = 30$), by supplier and objective weight λ (in %). CV measures between-tree variability of estimated expected outcomes. Environmental variability is consistently around 18 % to 20 % for $\lambda \leq 0.75$ and decreases to approximately 12 % to 13 % at $\lambda = 1$, while economic variability is comparatively stable across λ (approximately 9 % to 9.5 % for Suppliers A

and B, and around 11 % to 11.5 % for Supplier C). Overall, the seed-to-seed variation is materially larger for environmental outcomes than for economic outcomes, and Supplier C exhibits the largest relative variability on the economic dimension.

Table 5.11: Summary statistics for product-category contributions grouped by impact type.

		Descriptive statistics						
#	Category	μ	σ	min	$Q_{0.25}$	$Q_{0.50}$	$Q_{0.75}$	max
Environmental impact share [%]								
1	Conference chairs	6.14	5.65	0.18	1.38	4.74	9.66	34.77
2	Conference tables	1.33	1.25	0.21	0.60	0.91	1.53	13.22
3	Education chairs	9.44	3.92	2.52	6.56	8.81	11.67	54.58
4	Education tables	5.80	4.66	0.90	2.64	4.07	6.83	36.34
5	Electric office desks	1.62	1.73	0.13	0.64	0.88	2.01	15.39
6	Manual office desks	7.17	6.77	0.57	2.32	4.33	9.75	37.77
7	Office chairs	12.94	11.32	0.71	4.55	8.71	18.66	54.25
8	Pedestals	12.83	3.06	4.81	10.10	13.20	15.22	22.38
9	Storage cabinets	42.72	10.05	16.22	33.76	44.18	50.87	69.83
Economic impact share [%]								
1	Conference chairs	14.12	4.31	4.56	10.14	13.92	18.38	24.04
2	Conference tables	4.57	1.39	2.14	3.71	4.04	4.63	10.43
3	Education chairs	13.06	2.42	8.47	11.56	12.66	14.02	47.24
4	Education tables	5.84	3.16	2.36	3.76	4.86	6.45	26.19
5	Electric office desks	2.71	0.79	1.04	2.09	2.54	3.17	8.27
6	Manual office desks	18.51	1.95	10.09	17.18	18.47	19.79	27.38
7	Office chairs	18.31	3.34	8.38	15.97	17.61	19.94	30.26
8	Pedestals	6.83	0.74	3.35	6.33	6.81	7.31	10.39
9	Storage cabinets	16.05	1.95	7.67	14.70	15.98	17.32	23.45

Table 5.12: Coefficient of variation (CV) of the replication-level mean outcome estimates across scenario-tree seeds ($R = 30$), by supplier and objective weight λ (in %).

	Objective weight λ				
Supplier	0.00	0.25	0.50	0.75	1.00
Environmental Coefficient of Variation [%]					
A	18.45	18.42	18.55	18.34	11.91
B	17.93	17.72	17.66	17.64	11.80
C	19.05	19.70	19.74	19.74	13.29
Economic Coefficient of Variation [%]					
A	9.01	9.37	9.39	9.41	9.64
B	8.90	9.13	9.22	9.24	9.33
C	10.73	11.21	11.22	11.22	11.27

Table 5.13 summarizes the distribution of optimized process amounts that drive the impact decomposition; return flows are much less variable than forward flows. The table shows that manufacturing quantities are highly variable due to forward demand uncertainty, while return-flow quantities are comparatively stable, reflecting the smoothing effect of the aggregated return model.

Table 5.13: Summary statistics for optimized functional demand (forward and return). Entries report the scenario-contingent process amounts $S_{\lambda,r,s}$ (units: item flows) pooled over λ and scenario-tree replications r. Each observation corresponds to one (r, s) pair; hence n is identical across processes ($n = 43{,}740$). CV denotes σ/μ.

	Descriptive statistics							
Process	μ	σ	min	$Q_{0.25}$	$Q_{0.50}$	$Q_{0.75}$	max	CV
Forward processes								
Conference chair	409.35	25.58	328	392	407	425	1,346	0.062
Conference table	40.05	11.14	28	34	35	37	91	0.278
Education chair	410.68	148.66	279	328	368	453	4,212	0.362
Education table	206.97	98.06	127	151	157	189	749	0.474
Electric office desk	70.20	12.74	45	59	68	78	171	0.181
Manual office desk	395.05	27.29	317	375	391	413	582	0.069
Office chair	518.22	33.24	386	493	514	541	803	0.064
Pedestal	184.61	19.03	132	171	182	197	340	0.103
Storage cabinet	123.43	13.64	95	113	123	132	197	0.111
Return processes								
Conference chair (R)	44.36	1.28	38	44	44	45	52	0.029
Conference table (R)	5.94	0.24	4	6	6	6	8	0.040
Education chair (R)	378.05	3.30	362	376	378	380	401	0.009
Education table (R)	83.92	1.66	74	83	84	85	92	0.020
Electric office desk (R)	36.06	1.13	31	35	36	37	43	0.031
Manual office desk (R)	67.91	1.47	60	67	68	69	80	0.022
Office chair (R)	265.11	2.56	253	263	265	267	286	0.010
Pedestal (R)	38.12	1.29	33	37	38	39	48	0.034
Storage cabinet (R)	12.13	0.48	10	12	12	12	16	0.040

5.5.2 Variance Decomposition

Figure 5.12 demonstrates that the total variance of the aggregate impact $\tilde{H}(x; \lambda)$ is dominated by two distinct drivers, depending on the objective dimension. This section decomposes outcome uncertainty into (i) structural tree uncertainty (V_{tree}), (ii) scenario-choice uncertainty (V_{scen}), and (iii) process-data uncertainty (V_{proc}), following the decomposition defined in Equation (3.47).

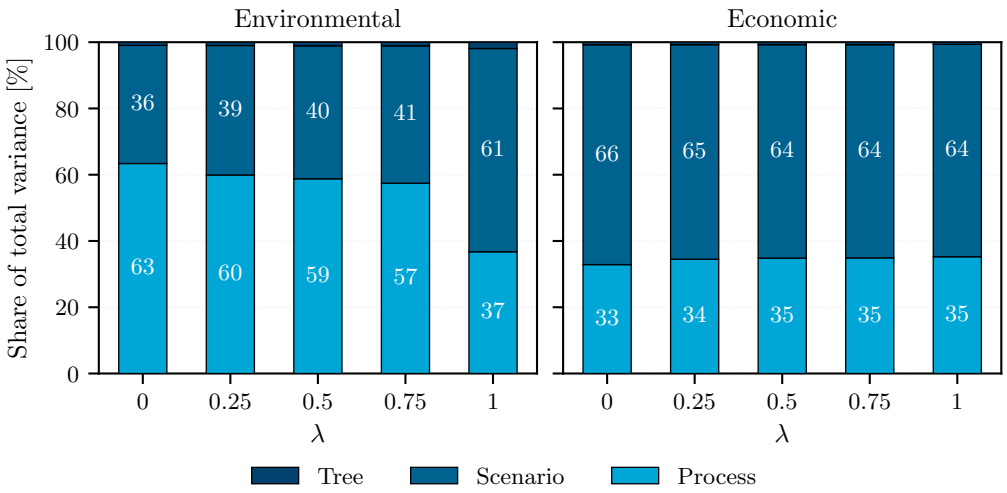

Figure 5.12: Share of the total variance of $\tilde{H}(x;\lambda)$ attributable to structural uncertainty (V_{tree}), scenario-choice uncertainty (V_{scen}), and process-data uncertainty (V_{proc}), shown across λ and averaged over suppliers.

Within the generated ensemble of scenario trees, structural uncertainty V_{tree} is consistently negligible ($< 1\%$ of total variance) across all λ. The dominant drivers split sharply between impact dimensions. For economic impacts, variance is driven almost entirely by scenario-choice uncertainty (V_{scen}). In contrast, for environmental impacts (at $\lambda \leq 0.75$), variance is dominated by process-data uncertainty (V_{proc}).

To dissect the epistemic component, Tables 5.14–5.17 decompose V_{proc} into contributions by product category and life cycle stage. Reported shares correspond to the analytic first-order contributions (C_i^{analytic}, Eq. (3.62)), consistent with the independent sampling of process impacts employed in the simulation (Section 3.8.3). Note that because impacts are sampled independently, these results isolate unit-process variances but do not capture cross-correlations inherent in the background database.

Table 5.14: Product-category variance shares (in % of V_{proc}) by supplier and objective weight λ (environmental impact).

Category	λ														
	0.00			0.25			0.50			0.75			1.00		
	A	B	C	A	B	C	A	B	C	A	B	C	A	B	C
Environmental variance share [%]															
Conference chairs	5.8	8.5	13.3	4.3	6.5	23.2	4.3	5.5	23.3	4.5	5.4	23.3	6.1	7.8	41.4
Conference tables	0.3	0.1	0.6	0.3	0.1	0.5	0.3	0.1	0.5	0.3	0.1	0.5	0.1	0.0	0.2
Education chairs	2.3	3.3	5.4	2.3	3.4	4.8	2.3	3.4	4.8	2.4	3.4	4.8	4.1	5.7	7.0
Education tables	2.1	1.5	1.0	2.3	1.6	0.8	2.6	1.8	0.8	2.7	1.9	0.8	7.4	5.7	0.8
Electric office desks	0.7	0.8	0.7	0.7	0.8	0.6	0.7	0.8	0.6	0.7	0.8	0.6	0.8	0.7	1.0
Manual office desks	16.0	12.6	10.9	16.1	12.9	9.6	15.9	13.0	9.6	15.5	12.5	9.6	3.2	12.4	2.1
Office chairs	68.1	68.3	63.7	69.2	69.8	56.5	69.1	70.4	56.4	68.9	70.8	56.4	69.7	59.2	41.7
Pedestals	0.3	0.3	0.4	0.3	0.3	0.4	0.3	0.3	0.4	0.3	0.3	0.4	0.5	0.5	0.6
Storage cabinets	4.4	4.6	4.0	4.5	4.7	3.6	4.5	4.8	3.6	4.7	4.8	3.6	8.0	8.0	5.2

Table 5.15: Product-category variance shares (in % of V_{proc}) by supplier and objective weight λ (economic impact).

Category	λ														
	0.00			0.25			0.50			0.75			1.00		
	A	B	C	A	B	C	A	B	C	A	B	C	A	B	C
Economic variance share [%]															
Conference chairs	62.1	65.8	60.7	66.5	68.6	63.8	66.5	69.8	63.9	66.5	69.9	63.9	72.2	71.4	66.1
Conference tables	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1	0.1
Education chairs	0.0	0.0	2.9	0.0	0.0	2.6	0.0	0.0	2.6	0.0	0.0	2.6	0.0	0.0	2.6
Education tables	0.2	0.1	1.2	0.3	0.1	1.6	0.3	0.1	1.6	0.3	0.1	1.6	0.3	0.1	1.6
Electric office desks	0.4	0.4	0.3	0.3	0.4	0.3	0.3	0.4	0.3	0.3	0.3	0.3	0.3	0.3	0.3
Manual office desks	1.7	1.5	1.4	1.5	1.4	1.3	1.5	1.4	1.3	1.5	1.4	1.3	1.3	1.3	1.3
Office chairs	35.4	32.2	33.4	31.2	29.4	30.3	31.2	28.3	30.3	31.2	28.2	30.3	25.9	26.7	28.0
Pedestals	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Storage cabinets	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

Table 5.16: Decomposition of the process-data variance component V_{proc} (environmental impact) into contributions by life cycle stage (rows) for each supplier and λ (columns). Values are percentages of V_{proc} and sum to 100% per column.

Life Cycle Phase	λ														
	0.00			0.25			0.50			0.75			1.00		
	A	B	C	A	B	C	A	B	C	A	B	C	A	B	C
Environmental variance share [%]															
Manufacturing	99.7	99.8	99.9	99.7	99.8	99.9	99.7	99.8	99.9	99.7	99.8	99.9	99.4	99.6	99.9
Warehousing	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Transportation	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Use phase	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
End-of-life	0.3	0.2	0.1	0.3	0.2	0.1	0.3	0.2	0.1	0.3	0.2	0.1	0.6	0.4	0.1

Table 5.17: Decomposition of the process-data variance component V_{proc} (economic impact) into contributions by life cycle stage (rows) for each supplier and λ (columns). Values are percentages of V_{proc} and sum to 100% per column.

Life Cycle Phase	λ														
	0.00			0.25			0.50			0.75			1.00		
	A	B	C	A	B	C	A	B	C	A	B	C	A	B	C
Economic variance share [%]															
Manufacturing	100	100	100	100	100	100	100	100	100	100	100	100	91.9	100	100
Warehousing	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Transportation	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	8.1	0.0	0.0
Use phase	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
End-of-life	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

An important contrast emerges when comparing variance contributions to mean impact shares. While Storage cabinets account for the largest share of mean environmental impact

($\approx$ 43%, see Table 5.11), they contribute only $\approx$ 4% to the variance. Conversely, Office chairs, which constitute only $\approx$ 13% of the mean impact, drive approximately 68% of the total process-data variance (Table 5.14). Choice of pathway (new vs. circular) is encoded in the activity vector and affects V_{scen}, while the parametric uncertainty of the chosen circular processes drives V_{proc}. The stage-wise decomposition (Table 5.16) confirms that process-data uncertainty is concentrated in the Manufacturing stage ($>$ 99%) for environmental impacts.

Chapter 6

Discussion

This chapter synthesizes the thesis methodological and empirical contributions. It provides (i) answers to the three research questions stated in Chapter 1 and (ii) an assessment of limitations and scope conditions that bound the interpretability and transferability of the results. The discussion closes by articulating theoretical, methodological, and practical implications and by identifying concrete directions for future research.

6.1 Summary of Findings

The primary contribution of this study is a unified methodological framework (RQ1) for supply chain life cycle optimization that jointly optimizes environmental and economic impacts under uncertainty. It integrates the algebraic structure of LCA within a multi-stage stochastic supply chain optimization model, making life cycle inventory and impact assessment endogenous and decision-dependent rather than an ex-post exercise. In this way, uncertainty in supplier-dependent life cycle variables (inventory uncertainty) and in exogenous drivers such as demand and returns (demand/return uncertainty) is treated consistently within a single coherent optimization framework.

Empirically (RQ2), the application to the TU Delft campus furniture case study demonstrates that the proposed framework can be operationalized in a real decision context where a substantial share of impacts is embedded in upstream supply-chain choices. The case study illustrates how heterogeneous suppliers and uncertain LCIA data, together with stochastic demand and return flows, can be structured into a tractable multi-period, capacity-constrained optimization model that links a long-term procurement commitment to adaptive recourse decisions with impact estimates. In this way, the framework provides a quantitative basis for sourcing and circularity planning under uncertainty, supporting near- to medium-term supply-chain decarbonization efforts.

In terms of evaluation insights (RQ3), the results indicate a trade-off between environmental and economic objectives, producing solution regimes that favor either GHG emission reduction or cost efficiency. The analysis of optimality gaps confirms that the efficient frontier is polarized: Supplier B outperforms in terms of economic impact, Supplier C minimizes environmental impact, while intermediate options (Supplier A) are consistently dominated. The uncertainty analysis further reveals that economic impact's sensitivity is primarily driven by scenario (functional demand) uncertainty (V_{scen}), whereas environmental impact is dominated by epistemic data uncertainty in the manufacturing stage (V_{proc}). Furthermore, the decision-value metrics highlight a distinction between strategic and operational robustness. The negligible VSS indicates that the strategic supplier choice is stable despite uncertainty, while the EVPI indicates that operational performance is sensitive to having perfect information about uncertain variables.

6.2 Interpretation of Results

6.2.1 Trade-offs and Supplier Performance

The results indicate a strongly asymmetric trade-off between environmental and economic performance. Figure 5.5 shows that moving from the environmentally preferred configuration (Supplier C, $\lambda = 0$) to the cost-minimizing configuration (Supplier B, $\lambda = 1$) reduces expected costs from roughly 612 k€ to 539 k€ (about 12 %), while increasing expected life cycle emissions from roughly 90 tCO_2eq to 182 tCO_2eq, doubling the life cycle GHG emissions. In the modeled setting, relatively small cost savings are therefore associated with large increases in emissions. Conversely, substantial emissions reductions appear achievable at comparatively modest additional cost, confirming that meaningful environmental improvements do not necessarily require large cost increases.

The shape of the Pareto-efficient set further indicates that these trade-offs are not constant along the frontier. Table 5.4 shows that over a broad intermediate range of preference weights ($\lambda \in [0.25, 0.75]$), expected emissions vary only slightly, whereas expected costs decrease by several tens of k€. As λ approaches the economic extreme, the pattern reverses: further reductions in expected costs become marginal, while emissions increase sharply. This implies a changing marginal rate of substitution along the frontier: early moves away from the environmental extreme mainly buy cost reductions at limited environmental penalty, whereas later moves mainly degrade environmental performance for little additional economic gain. This relationship emphasizes the importance of multi-criteria analysis in SCLCO.

Across the efficient set, the frontier is largely determined by solutions associated with Suppliers B and C. Supplier A is not competitive when both objectives are considered jointly and is frequently dominated in the bi-objective space. This aligns with the optimality-gap results in Section 5.3. Under pure environmental minimization ($\lambda = 0$), Supplier C achieves performance statistically indistinguishable from the best SAA benchmark, whereas under pure cost minimization ($\lambda = 1$), Supplier B is effectively optimal.

This structure persists for intermediate preference weights ($\lambda \in \{0.25, 0.50, 0.75\}$). Supplier C remains near-optimal at low weights but becomes clearly suboptimal as the economic objective is emphasized, while Supplier B becomes near-optimal once $\lambda \geq 0.50$. Supplier A exhibits consistently positive optimality gaps throughout the intermediate regime, indicating that it is seldom selected by the unrestricted SAA benchmark and that choosing Supplier A typically entails an avoidable loss across a wide range of preference weights. Overall, the strategic choice is primarily between Supplier C (superior environmental performance) and Supplier B (lowest expected cost).

The step-change diagnostics provide a behavioral interpretation of how portfolios move along the frontier. Focusing on the production–circularity decision (new production versus refurbishment), Table 5.5 shows a consistent reallocation away from refurbishment and toward new production as λ increases. The direction of this shift is already visible between $\lambda = 0.00$ and 0.25 (new production increases and refurbishment decreases for all suppliers), but the largest adjustment occurs in the final transition from $\lambda = 0.75$ to 1.00, where new production rises sharply and refurbishment falls comparably across suppliers. This indicates that the economic extreme is associated with a discrete move toward new sourcing rather than a smooth continuation of the mixed-objective regime, which is consistent with the steep environmental penalty observed near $\lambda = 1$.

The exceedance diagnostics further indicate that these portfolio shifts are not primarily driven by binding circularity shortfalls. As shown in Table 5.6, where defined, circularity exceedances remain below 10%, implying that available capacities cover the large majority of demand and that the solutions require little λ-invariant forced hedging to satisfy feasibility.

From a decision-support perspective, the key issue becomes where to locate the decision boundary along the efficient frontier: the region in which meaningful emissions reductions can be secured at an acceptable additional cost, versus the region in which additional cost savings are small and associated with rapidly increasing environmental penalties.

6.2.2 The Value of Information Modeling

The value-of-information results show a contrast between the EVPI and the VSS. A large EVPI (e.g., ≈21% of emissions at $\lambda = 1$) combined with a negligible VSS implies a specific decision structure. The near-zero VSS indicates that the first-stage supplier selection is stable: for a given λ, the same supplier is optimal in the expected-value model and in the stochastic recourse model. This is consistent with suppliers occupying well-separated regions of the environmental–economic trade-off space, such that the modeled uncertainty does not shift the optimal supplier choice.

This result should also be interpreted in light of the scenario-approximation diagnostics. The scenario trees reproduce mean demand well but understate dispersion, which may under-represent tail events. In a risk-neutral setting, VSS reflects improvements in expected value only; if extreme realizations where hedging would matter are not well represented, the stochastic model has limited incentive to change first-stage decisions relative to the expected-value plan. Under this interpretation, the small VSS reflects both decision robustness and the limits of the discretized uncertainty set in capturing tail-driven benefits of conservative planning. For the evaluated instance, explicitly modeling uncertainty therefore yields little improvement in expected performance relative to a expected-value approximation.

In contrast, the high EVPI indicates value in foresight. Because the recourse problem is non-anticipative, it must hedge across demand and return paths. Under wait-and-see information, supplier choice and subsequent inventory and flow decisions can be conditioned on the realized trajectory. The EVPI thus quantifies the cost of acting under partial information through inventory positioning and production/recovery plans, even when the ex ante supplier choice is stable.

Finally, even when VSS is close to zero, stochastic modeling remains useful for characterizing outcome variability. Representing plausible scenarios provides insight into the range and likelihood of total emissions over the planning horizon, rather than only an expected-value point estimate.

6.2.3 Uncertainty and Sensitivity Analysis

The contribution analysis (Section 5.5.1) shows a consistent contrast across suppliers: Supplier C has lower environmental burdens but higher (and more variable) economic outcomes, while Supplier B is economically favorable and environmentally similar to Supplier A.

The variance decomposition (Section 5.5.2) further shows that the main sources of uncertainty differ by impact dimension. Which scenario is realized (V_{scen}) drives most of the variability in economic output (exceeding process-data uncertainty), reflecting exogenously specified, scenario-specific differences in demand and return volumes that the optimized system must satisfy. Environmental variance (for $\lambda \leq 0.75$), in contrast, is dominated by process-data uncertainty (V_{proc}). Environmental uncertainty is therefore mainly epistemic: it arises from uncertain inventory parameters, including ecoinvent-based background data and probabilistic elements in the foreground circularity model.

Comparing mean impact shares with variance contributions highlights a clear mismatch between what drives expected impacts and what drives uncertainty. Storage cabinets account for the largest share of mean environmental impact ($\approx 43\%$) but contribute only $\approx 4\%$ of the variance. Office chairs, by contrast, make up only $\approx 13\%$ of the mean impact yet account for about 68% of the total process-data variance. This pattern follows directly from the modeling structure. Office chairs include several probabilistic circular pathways (e.g., maintenance, refurbishment, revitalization) and multiple component-level revitalization choices, such as reupholstery or replacement of the seat/padding, gas spring, and castors, whereas storage cabinets are modeled as non-circular assets. The decomposition therefore suggests that the added uncertainty introduced by these pathway and component-level choices dominates the process-data variance. This effect is layered on top of the background inventory uncertainty shared across products (including raw material production) and reflects sensitivity to refurbishment and revitalization (hyper)parameters.

The results per life cycle phase support this interpretation. For environmental impacts, process-data uncertainty is concentrated in the manufacturing stage (> 99%). Transport and end-of-life may contribute to expected impacts, but their process-data uncertainty is small relative to manufacturing and refurbishment inventories. For decision support, this distinction matters because it points to different levers for uncertainty reduction: when V_{proc} dominates, the most effective actions are to improve process data and constrain circularity parameters; when V_{scen} dominates, uncertainty is driven by volume variability and is less responsive to process-level refinements.

6.3 Methodological Reflection and Limitations

6.3.1 Validity of the Scenario Approximation

The validation of the scenario-tree approximation (Section 5.1) reveals limitations that shape how the results should be interpreted.

First, the reduction of life cycle impact coefficients preserves the underlying distributions well for environmental parameters: the portfolio-level median range-normalized distance is about 3 % of the outcome range (Table 5.1). Economic parameters are captured slightly less accurately, with a higher median normalized distance of about 5.5 %–6 % (Table 5.1). The economic discrepancies also exhibit heavier tails and greater sensitivity to the selected realizations (Figure 5.1). This indicates that expected environmental outcomes are estimated more robustly, whereas economic tail risk may be somewhat underrepresented due to the more discretized nature of the underlying distribution.

Second, while the scenario generation reproduces mean demand levels, the tree-implied standard deviations remain consistently below the empirical standard deviations (Figure 5.3). The scenario trees therefore represent typical operating conditions more faithfully than extreme realizations. Results that depend on volatility, e.g., service levels under high-demand or low-return scenarios, should be interpreted with caution. This under-dispersion also indicates that the scenario-generation procedure would benefit from refinement, for instance by placing greater emphasis on matching second moments or otherwise improving tail representation.

Finally, using 30 independent seeds reduces sampling noise and yields stable mean estimates, but repeated sampling cannot correct structural approximation errors induced by a finite scenario tree. In other words, replication increases precision around the tree-based approximation, yet it cannot recover empirical accuracy or missing variability.

Overall, the diagnostics suggest the scenario-tree approximation is more reliable for conclusions based on expected performance than for tail-dependent metrics. Because dispersion is systematically understated, results that hinge on extreme realizations (e.g., worst-case environmental or demand outcomes) should be interpreted conservatively.

6.3.2 Data Quality, Reliability, and Baseline Assumptions

Data availability and quality remain a matter of concern. The reliability model relies on Weibull distribution (hyper)parameters (shape and scale) that govern the simulated generation of return flows. These parameters were not estimated from campus-specific return histories but derived from engineering assumptions. While the variance decomposition identifies process-data uncertainty (V_{proc}) rather than volume volatility as the dominant driver of environmental risk, economic outcomes remain sensitive to scenario variability (V_{scen}) induced by these assumed reliability inputs. The assumed failure rates and aging profiles simplify complex degradation processes and may not fully capture early-life failures, heterogeneous usage patterns, or the influence of maintenance intensity on product longevity.

Similarly, the demand scenario generation is constrained by sparse historical order records. While sufficient to estimate aggregate intermittency parameters, the limited and sparse history reduces the ability to identify specific trends and multi-year seasonality, which may lead to an underestimation of low-frequency demand variability.

The analysis is also limited by the granularity and alignment of supplier disclosures. The product portfolio was modeled using representative public disclosures on BoMs, yet some mismatch between historical demand, university information systems, and available supplier disclosures was unavoidable. Variant-specific demand within product categories was therefore treated conservatively as homogeneous, which may mask niche variants with disproportionate impacts or costs and could shift the efficient frontier at a more granular level.

Uncertainty characterization relies on the pedigree matrix approach to construct uncertainty distributions for life cycle inventory inputs. This provides a standardized, literature-consistent method for translating qualitative data indicators into quantitative dispersion assumptions, but it imposes interpretative boundaries. The resulting distributions primarily reflect epistemic (data-quality) uncertainty, limited precision and reliability of available inventory data, rather than empirical process variability (aleatory uncertainty). Moreover, the translation of pedigree scores into geometric standard deviations inherits subjective judgment and employs generic scaling factors that may not match the variance structure of the furniture supply chain or supplier-specific differences. Consequently, the uncertainty propagated to environmental outcomes should be interpreted as decision risk arising from imperfect information, not purely as physical outcome variability. Reducing this uncertainty therefore requires targeted data improvement, for instance replacing proxy inventories with high-quality primary measurements and more granular disclosures from suppliers.

In addition, several key engineering and technology parameters (e.g., recycling yields, recycled content shares, and Weibull shape parameters) were treated as fixed baseline assumptions. The scenario-based optimization therefore propagates uncertainty conditional on these baseline values. As a result, the study can localize uncertainty (e.g., by identifying which life cycle stages dominate variance), but it does not quantify which of these baseline assumptions are most influential in driving that uncertainty, nor how optimal policies would change under alternative parameterizations. A systematic sensitivity analysis with re-optimization under perturbed parameter assumptions was beyond the scope of this case study; the reported trade-offs should therefore be interpreted as conditional on the specified technological and reliability characteristics of the portfolio.

6.3.3 Structural and Operational Model Scope

The optimization framework employs a flow-based representation of the supply chain. While computationally tractable, this formulation has structural limitations: it does not explicitly represent the state, making it infeasible to track time-evolving states of individual assets. The model also does not allow demand backlogging; all functional demand must be met within the decision stage. Consequently, strategies that rely on delayed fulfillment to smooth capacity peaks cannot be evaluated, and service level agreement (SLA) violations are not explicitly penalized beyond the costs associated with required recourse actions.

Additional limitations arise from the case-study context and the relatively passive nature of institutional procurement. The decision-maker (the university) primarily selects flows rather than actively operating industrial assets. As a result, the availability and performance of technologies and methods are largely market-determined, and their integration into this framework is limited to the information requested and disclosed. This reduces the operational leverage available to mitigate risk compared to an integrated manufacturer–operator setting.

Furthermore, the current scope excludes social impacts and focuses on global warming potential (GWP) as the sole environmental indicator. While common in climate-focused studies, this omits trade-offs with other impact categories such as resource depletion or toxicity.

Finally, the planning horizon was limited to six months to keep the multistage stochastic mixed-integer program computationally tractable. While sufficient to demonstrate the methodological mechanics of the SCLCO framework, the abbreviated horizon does not

capture full annual demand patterns (i.e., July–December) in the case study and may understate the long-run benefits of circular strategies (e.g., repeated revitalization cycles) and multi-year supplier constraints.

6.4 Case Study Implications

The TU Delft furniture case study demonstrates that supply chain life cycle optimization can be used as a decision-support layer for university procurement: it can translate Scope 3 objectives into supplier rankings, quantify the cost of environmental ambition, and specify recourse actions (e.g., recovery and replacement flows) that remain feasible under uncertain demand and returns. The practical value is therefore not only to select the optimal supplier, but to structure procurement so that environmental and economic performance is measurable, comparable, and enforceable over the contract horizon.

A first implication for future tenders is how to operationalize the environmental–economic trade-off. Rather than treating sustainability as a qualitative criterion, the university can use Pareto frontiers to set an explicit willingness-to-pay for emissions reductions (or, equivalently, an emissions budget for a given expenditure envelope). This supports procurement decisions that are transparent and defensible: the university can justify selecting a higher-cost supplier when the implied emissions reduction is large relative to the cost premium, and it can avoid supplier configurations where additional cost savings are marginal relative to the associated increase in Scope 3 emissions.

Second, the results indicate that the most effective levers for improving the reliability of procurement outcomes differ by objective. For economic robustness, the priority is reducing volume uncertainty: improved forecasting, better internal tracking of asset usage and returns, and contract clauses that buffer demand swings (e.g., volume bands, flexible call-off, or option-like capacity reservations) are likely to yield the largest benefit. For environmental robustness, the priority is reducing upstream epistemic uncertainty. This suggests that future procurements should place greater emphasis on data quality and disclosure obligations, for instance by requiring supplier- and component-specific LCI documentation, verification of recycled content and recycling yields, and transparency on manufacturing inventories for high-variance product categories. In short, better primary data and standardized disclosures are more valuable for reliable Scope 3 management than further refinement of consumption forecasts.

Third, the case study suggests that universities should evaluate offers under stress as well as under typical operating conditions. Because operational impact is sensitive to realized demand and return trajectories, supplier assessment should include scenario-based stress tests (e.g., high-demand/low-return periods) and should reward configurations that maintain service performance with bounded recourse costs. Embedding such scenario testing in the evaluation procedure makes procurement more resilient and reduces the likelihood that a supplier and its estimated impacts are selected based on average performance while performing poorly under adverse realizations.

6.5 Future Work

This thesis should be viewed as a starting point for stochastic supply chain life cycle optimization. The proposed SCLCO framework is intentionally scoped to a tractable set of indicators, uncertainty representations, and planning horizons. Several extensions would strengthen the framework methodologically and broaden its decision relevance.

A first and immediate extension is to move beyond GWP and incorporate additional LCIA categories (e.g., resource depletion, toxicity-related indicators, particulate matter formation, and water use). Expanding the impact vector would enable explicit analysis of cross-impact trade-offs and potential burden shifting, rather than implicitly assuming that climate-optimal policies are globally sustainable. Relatedly, the framework could be extended to include social sustainability, which is currently under-addressed in life cycle studies. Incorporating social LCA indicators and institutional constraints (e.g., labor

standards, living-wage requirements, local sourcing mandates, and auditability criteria) would enable a more holistic sustainability assessment that also better reflect public procurement objectives.

The current evaluation emphasizes expected-value performance and post-optimality uncertainty assessment. Future work could integrate risk attitudes directly into the optimization objective and constraints. Concretely, tail-sensitive formulations such as conditional value at risk (CVaR), chance constraints, or utility-based objectives could be used to explicitly hedge against high-impact, low-probability realizations that are underrepresented in risk-neutral objectives, especially relevant given the observed under-dispersion of the scenario trees. Such extensions would also allow robustness to be defined in terms of service reliability and worst-case environmental performance, not only in terms of mean outcomes.

A structural limitation of the current model is its flow-based representation, which cannot track individual assets and their evolving condition. Moving to a state-based (or hybrid) formulation would enable condition-dependent decisions and explicit maintenance histories, while also anchoring the model in a more established state-space literature that is directly compatible with standard stochastic control and approximate dynamic programming techniques.

The thesis conducts sensitivity and uncertainty analysis largely ex post. A valuable extension is to analyze policy stability endogenously: how much perturbations in inputs (e.g., LCI coefficients, reliability parameters, or demand/return processes) change the optimal first-stage commitments and recourse policies. This can be addressed via structured re-optimization under input perturbations, distributionally robust optimization, or ambiguity sets over uncertain parameters and scenario probabilities. In addition, re-optimization under alternative structural hypotheses (e.g., alternative circularity pathways, different end-of-life allocation rules, or different representations of supplier constraints) would clarify which insights are robust to modeling choices rather than merely to parameter noise.

The six-month horizon was chosen for computational tractability, but it truncates annual cycles and multi-year asset lifetimes. Extending the time horizon to full-year and multi-year settings would enable evaluation of repeated refurbishment loops, long-term supplier commitments, learning effects, and capacity evolution. Practically, this can be coupled with rolling-horizon or model-predictive control approaches, where policies are re-optimized as new information becomes available and forecasts update. Such a design would reflect institutional procurement practice more closely and would allow explicit valuation of recourse flexibility under sequential information revelation.

Several key inputs (e.g., Weibull parameters governing returns and failure dynamics, refurbishment yields, and pathway probabilities) are currently based on engineering assumptions. Future work should leverage longitudinal asset-tracking and maintenance data to estimate failure, return, and refurbishment probabilities empirically, reducing reliance on assumed reliability profiles and improving realism of circular flow modeling. Integrating reliability engineering more explicitly, including early-life failures, heterogeneous usage, and maintenance-dependent hazard rates, would likely improve both predictive accuracy and the credibility of decision support.

The results indicate that environmental uncertainty is dominated by process-data uncertainty in manufacturing, which points to the value of improved primary data. An important direction is therefore to study how access to higher-resolution, supplier-verified LCI data alters optimal solutions, uncertainty attribution, and supplier rankings relative to background-database-based inventories. This is nontrivial in practice, since many current EPDs still rely on generic background data (e.g., ecoinvent) for upstream processes. However, certain domains, notably the electricity and energy sectors (e.g., origin-specific grid mixes, wind/photovoltaic sourcing, contractual instruments), may provide settings where higher-fidelity origin data are available and where the framework can appropriately integrate this information.

The current case study reflects a single decision-maker selecting among exogenously specified suppliers. Extending the framework to settings with multiple buyers or strategic

suppliers would allow modeling of competition, information asymmetry, and endogenous supplier responses (e.g., suppliers adjusting prices, disclosure, or capacity in response to procurement criteria). Such extensions could draw on bilevel optimization, equilibrium models, or mechanism-design-inspired formulations to represent how procurement rules shape supplier incentives and, in turn, life cycle performance.

Finally, the framework is applicable to many more product categories, organizational contexts, and geographic settings to identify scope conditions and transferable patterns. Applying the model to other circular assets (e.g., IT equipment, lab devices, building components) and other public and private organizations environments can assist in future development of the framework and its purpose.

6.6 Conclusion

This thesis addressed a central difficulty in organizational decarbonization: decisions that affect a large share of life cycle impacts must be taken sequentially, under imperfect information, and with economic constraints. It responds by a methodology that embeds the algebraic structure of LCA within a multistage stochastic supply chain optimization model, so that environmental and economic impacts become decision-dependent quantities that shape both long-term commitments and adaptive recourse actions.

The broader implication is methodological as much as practical. Integrating life cycle uncertainty into the supply chain optimization problem changes what can be concluded from an analysis: it produces policies rather than point estimates, makes trade-offs explicit rather than implicit, and allows uncertainty to be traced to the full life cycle components that generate it. As such, the framework provides a concrete basis for evaluating and comparing alternative supply-chain configurations when both impacts and operating conditions are uncertain, and it clarifies what additional data and modeling fidelity are most valuable for improving decision quality.

The quantitative conclusions remain bounded by the case-study scope and by the quality of available data and scenario approximations, but the framework provides a concrete basis for turning decarbonization targets into actionable decisions and for prioritizing the next improvements in data and modeling.

Appendix A

Case-study Details

Table A.1: Flow stages, product categories, and process types represented in the optimization model. Items labeled as *broken* denote units requiring reverse-flow handling (e.g., due to damage or return criteria), not necessarily physical breakage.

Model element	**Categories**
Flow stages	Produced; Collected; Warehoused; Delivered; Installed; Returned (broken); Collected (broken); Warehoused (broken); Delivered (broken)
Products	Office chair; Manual office desk; Electric office desk; Conference chair; Conference table; Education chair; Education table; Pedestal; Storage cabinet
Processes	New production; Refurbished production; Broken item revitalization; Broken item disposal; Warehousing; Collection; Transportation (diesel); Transportation (electric); Installation

Table A.2: Conditional probabilities of component failures given a product is classified as broken. Sampling proceeds until the first failed component is realized; failure is guaranteed by construction, so every return triggers either partial repair or full replacement.

Product group	Failure description	Probability
Office chairs	Armrests repair	0.150
Office chairs	Casters (set of five) repair	0.400
Office chairs	Gas spring repair	0.250
Office chairs	Seating repair	0.120
Office chairs	Upholstery repair	0.070
Manual office desks	Crank repair	0.456
Manual office desks	Tabletop repair	0.452
Electric office desks	Electric parts repair	0.070
Electric office desks	Tabletop repair	0.452
Education tables	Tabletop repair	0.540
Conference tables	Tabletop repair	0.540
Conference chairs	Seating repair	0.267
Conference chairs	Upholstery repair	0.144

Note: The probabilities shown are modeling assumptions used to parameterize conditional repair-category selection given that an item is classified as broken. They should not be interpreted as observed frequencies or statistically estimated component-failure probabilities. Because component-level failure microdata were unavailable, the parameter magnitudes are not empirically calibrated. Sources (only) motivate the conceptual modeling choice: the repair-category taxonomy support and its qualitative prioritization (including the salience of mobility/adjustment mechanisms for chairs and the non-independence of upholstery condition and seating support/padding in chair assessment) (Dorland & Jørgensen, 2025; Heller-Ono, 2020; Jetti & Dhar, 2024).

Table A.3: Probability that a broken item is classified as non-repairable (scrap), implying full replacement. The probabilities shown are modeling assumptions used to parameterize the conditional selection of repair categories given that an item is classified as broken.

Product group	Classification	Probability
Office chairs	Non-repairable (scrap)	0.100
Manual office desks	Non-repairable (scrap)	0.050
Electric office desks	Non-repairable (scrap)	0.050
Education tables	Non-repairable (scrap)	0.050
Conference tables	Non-repairable (scrap)	0.050
Conference chairs	Non-repairable (scrap)	0.100

Table A.4: Product variants included in the case study.

Product category	**Description**
Conference chairs	Conference chair with chrome sled base and integrated armrests; upholstered seat and back; arm pads.
Conference chairs	Conference chair with mobile pedestal base and armrests; upholstered seat and back; suitable for meeting and collaboration spaces requiring mobility.
Conference chairs	Four-leg conference chair, stackable, without armrests; upholstered seat pad; fitted with hard glides for smooth indoor flooring.
Conference tables[b]	Conference table on a steel frame with outward-set legs; fixed-height meeting table format.
Conference tables[b]	Conference table with melamine worksurface on a steel frame with square legs; fixed-height meeting table format.
Conference tables[b]	Round meeting table with cross base; fixed-height meeting table format.
Education chairs	Shell chair with five-star base; swivel function; soft casters; gas-lift height adjustment; plastic shell; shell available in multiple colors.
Education chairs	Four-leg shell chair, stackable, without armrests; plastic seat without upholstery; equipped with glides; shell available in multiple colors.
Education tables[b]	Education table with melamine worksurface on a steel frame with square legs; fixed-height classroom table format.
Electric office desks[a]	Electric sit–stand desk on an I-leg frame; height-adjustable workstation desk.
Electric office desks[a]	Duo electric sit–stand workstation (two workstations on a shared frame, no divider); otherwise equivalent in function and construction to the single sit–stand desk.
Manual office desks[a]	Crank height-adjustable desk on an I-leg frame; manually height-adjustable workstation desk.
Manual office desks[a]	Duo crank-adjustable workstation (two workstations on a shared frame, no divider); otherwise equivalent in function and construction to the single crank-adjustable desk.
Office chairs	Task chair with adjustable lumbar support and synchro-tilt mechanism; seat-depth adjustment; 3D adjustable armrests; casters suitable for hard and soft floors; upholstered seat; seat height adjustable.
Pedestals	Mobile pedestal with lock; three drawers including a materials drawer.
Storage cabinets	Low sliding-door storage cabinet with lock; supplied with internal shelves; perforated panels providing acoustic absorption.

[a] **Desk normalization:** Duo workstations are converted to two equivalent solo desks for analysis. All desks are treated as having a standardized footprint of 160×80 cm.
[b] **Table normalization:** Education and conference tables are converted to 200×100 cm tabletop-area equivalents.

Table A.5: Overview of material, electronic, chemical, composite, mineral, and production-related inputs included in the case study as raw materials.

Category	Inputs included
Metals	Aluminium; steel (low-alloyed, unalloyed); Zamak
Polymers	ABS; EVA; LDPE; PE; PA6; PA66; polyester; PET; POM; PP; PS; PU (foam, plastic); rubber
Wood	Chipboard; multiplex plywood; oak
Natural fibres	Metisse; wool
Minerals / construction materials	Concrete
Electronics	Cables; control units; motors; switches and displays; electrification parts
Chemicals and auxiliaries	Epoxy; glue; grease; lacquer; powder coating; melamine
Composites	Melamine paper
Production-related inputs	Steel welding

Abbreviations: ABS = acrylonitrile–butadiene–styrene; EVA = ethylene–vinyl acetate; PE = polyethylene; LDPE = low-density polyethylene; PA = polyamide; PET = polyethylene terephthalate; POM = polyoxymethylene; PP = polypropylene; PS = polystyrene; PU = polyurethane.

Table A.6: Recycling-rate and recycled-content assumptions used in the material modeling (non-zero parameters shown).

Category	**Material**	**Parameter**	**Value**
Metals	Aluminium	Recycled content[b]	32 %
	Aluminium	Recycling rate[a]	85 %
	Low-alloy steel	Recycled content[b]	35.5 %
	Low-alloy steel	Recycling rate[a]	90 %
	Unalloyed steel	Recycled content[b]	35.5 %
	Unalloyed steel	Recycling rate[a]	90 %
	Zamak	Recycling rate[a]	32 %
Polymers	LDPE	Recycling rate[a]	35 %
	PA6	Recycling rate[a]	10 %
	PA66	Recycling rate[a]	60 %
	Polyester	Recycling rate[a]	37.5 %
	Polyethylene (PE)	Recycling rate[a]	35 %
	PET	Recycling rate[a]	64.5 %
	POM	Recycling rate[a]	25 %
	PP	Recycling rate[a]	20 %
	Polystyrene (PS)	Recycling rate[a]	43 %
	PU foam	Recycling rate[a]	40 %
	PU plastic	Recycling rate[a]	5 %
Wood	Chipboard	Density	700 $\mathrm{kg\,m^{-3}}$
	Chipboard	Recycling rate[a]	33 %
	Multiplex plywood	Density	650 $\mathrm{kg\,m^{-3}}$
	Oak	Density	640 $\mathrm{kg\,m^{-3}}$
Other	Concrete	Recycling rate[a]	70 %
	Control unit (electronics)	Recycling rate[a]	40.6 %
	Motor (electronics)	Recycling rate[a]	40.6 %
	Cable (electronics)	Recycling rate[a]	40.6 %
	Electrification parts	Recycling rate[a]	40.6 %
	Powder coating	Renewable electricity share	100 %

Default assumption: Parameters not listed in this table are assumed to be zero.
[a] **Recycling rate:** End-of-life share of material that is collected and processed for recycling.
[b] **Recycled content:** Inflow share of recycled material used in production; may include **PCR** (post-consumer recycled) and/or **PIR** (post-industrial recycled).

Table A.7: The cohort-based estimated mean monthly returns by product category over 2025.

Product category	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Annual
Office chairs	42.0	42.9	43.7	44.6	45.5	46.5	47.4	48.3	49.1	50.1	51.0	51.9	563.0
Manual office desks	10.6	10.9	11.1	11.4	11.7	12.0	12.3	12.6	12.9	13.2	13.5	13.8	145.9
Electric office desks	5.6	5.8	5.9	6.0	6.1	6.3	6.4	6.6	6.7	6.8	7.0	7.1	76.2
Education chairs	60.8	61.8	62.6	63.4	64.3	65.1	65.9	66.7	67.5	68.3	69.0	69.8	785.2
Education tables	13.3	13.6	13.9	14.2	14.5	14.8	15.1	15.4	15.7	16.0	16.3	16.6	179.4
Conference chairs	6.8	7.0	7.3	7.5	7.7	7.9	8.2	8.4	8.6	8.9	9.1	9.4	96.9
Conference tables	0.5	0.5	0.5	0.6	0.6	0.6	0.6	0.6	0.7	0.7	0.7	0.7	7.3
Pedestals	6.0	6.1	6.3	6.4	6.6	6.8	6.9	7.1	7.3	7.4	7.6	7.8	82.4
Storage cabinets	1.7	1.7	1.8	1.8	1.9	2.0	2.0	2.1	2.1	2.2	2.2	2.3	23.8
Total	147.4	150.3	153.1	156.0	158.9	161.9	164.8	167.8	170.5	173.5	176.4	179.4	1,960.0

Table A.8: Key parameter assumptions for heavy-duty road transport modelling.

Assumption	Value
Functional unit	
Distance basis Functional unit used for transport modelling.	1 t km
Average payload Average transported mass per trip.	5.8 t
Vehicle lifetime Lifetime driving distance of a heavy-duty lorry.	540,000 km
Common structure	
Base dataset Background process used for heavy-duty transport modelling.	ecoinvent 3.11, lorry 16–32 t, EURO 6 (RER)
Foreground scope Foreground elements included in the transport model.	Vehicle production, maintenance, tyres, road wear
LCIA method Climate-change impact assessment method.[a]	IPCC 2021 GWP_{100}
Diesel lorry	
Process name Custom process name used for diesel transport modelling.	Heavy-duty lorry transport, diesel EURO 6, 16–32 t
Fuel and exhaust Use-phase fuel consumption and tailpipe emissions.	As in ecoinvent (unmodified)
Brake, tyre and road wear Non-exhaust emissions associated with vehicle operation.	As in ecoinvent (unmodified)
Battery-electric lorry	
Process name Custom process name used for battery-electric transport modelling.	Heavy-duty lorry transport, BEV, 16–32 t
Electricity use Electricity consumption per unit transport service.[b]	0.10 kW h t^{-1} km
Battery pack size Installed battery energy capacity.	540 kW h
Battery climate impact Manufacturing emissions per battery pack.	34,000 kg CO_2−eq
Battery lifetime allocation Battery manufacturing impacts allocated per unit transport service.	$1/(540{,}000\,\text{km} \times 5.8\,\text{t})$
Brake wear factor Brake-wear reduction factor relative to diesel lorry.	50 %

[a] Excluding biogenic CO_2.
[b] Netherlands medium-voltage grid mix.

Table A.9: Modeling assumptions used for the warehousing processes. Values are taken directly from the model configuration and apply uniformly across all warehousing activities.

Assumption	Value
Infrastructure	
Building lifetime Service life of warehouse building used for infrastructure allocation.	50 year
Building dataset region Region assumed for warehouse construction (steel hall).	CH
Functional unit Basis for warehousing service provision.	1 m^2 month
Energy use	
Electricity intensity (annual) Electricity use per unit floor area per year.	33 kW h m^{-2} $year^{-1}$
Heat intensity (annual) Natural-gas heat use per unit floor area per year.	47 kW h m^{-2} $year^{-1}$
Electricity per m^2 month Monthly electricity demand derived from annual value.	2.75 kW h
Heat per m^2 month Monthly heat demand derived from annual value and converted to MJ.	14.10 MJ
Geography	
Warehousing location Assumed geographic location of warehousing operations.	NL
Electricity dataset region Electricity grid mix used for warehouse electricity consumption.	NL
Heat dataset region Heat production technology used for warehouse space heating.	RER
Area densities	
Office chair Items stored per square meter of warehouse floor area.	20.83
Desk (160×80) Items stored per square meter of warehouse floor area.	15.63
Table (160×80) Items stored per square meter of warehouse floor area.	15.63
Table (200×100) Items stored per square meter of warehouse floor area.	10.00
Storage cabinet Items stored per square meter of warehouse floor area.	20.83
Visitor chair Items stored per square meter of warehouse floor area.	20.83
Stackable visitor chair Items stored per square meter of warehouse floor area.	83.33
Pedestal unit Items stored per square meter of warehouse floor area.	31.25

Appendix B

Additional Results

Table B.1: Summary statistics for environmental and economic impacts by supplier, reported separately for each λ (no pooling).

		Descriptive statistics							
Supplier	n	μ	σ	min	$Q_{0.25}$	$Q_{0.50}$	$Q_{0.75}$	max	CV
$\lambda = 0.00$									
Environmental impact [tCO_2eq]									
A	10,000	121.30	26.77	75.18	102.99	114.08	132.08	262.80	0.221
B	10,000	123.41	26.50	72.02	104.76	117.15	134.74	244.85	0.215
C	10,000	93.75	21.18	56.80	78.83	88.09	103.18	225.00	0.226
Economic impact [k€]									
A	10,000	595.45	53.96	456.62	557.30	591.82	628.67	914.80	0.091
B	10,000	567.21	50.86	442.35	530.86	562.29	599.09	875.32	0.090
C	10,000	616.23	66.47	455.34	568.56	606.48	654.63	1,075.43	0.108
$\lambda = 0.25$									
Environmental impact [tCO_2eq]									
A	10,000	124.89	27.35	74.72	106.19	118.26	135.76	279.43	0.219
B	10,000	126.34	26.70	72.68	108.01	120.27	137.68	246.25	0.211
C	10,000	98.44	22.39	56.73	82.06	93.52	109.29	225.79	0.227
Economic impact [k€]									
A	10,000	589.74	55.83	444.56	550.62	586.19	624.15	913.54	0.095
B	10,000	563.61	52.02	432.87	526.26	558.75	596.76	875.14	0.092
C	10,000	610.30	68.73	444.82	561.25	600.65	649.50	1,075.43	0.113
$\lambda = 0.50$									
Environmental impact [tCO_2eq]									
A	10,000	125.49	27.59	74.69	106.53	118.83	137.03	279.60	0.220
B	10,000	127.78	26.86	72.67	109.24	121.59	139.41	247.15	0.210
C	10,000	98.51	22.45	56.73	82.05	93.57	109.48	225.89	0.228
Economic impact [k€]									
A	10,000	588.83	55.84	443.85	549.59	585.34	623.25	912.07	0.095
B	10,000	562.40	52.40	432.62	524.75	557.62	595.76	874.85	0.093
C	10,000	610.28	68.73	444.82	561.23	600.63	649.50	1,075.43	0.113
$\lambda = 0.75$									
Environmental impact [tCO_2eq]									
A	10,000	127.60	27.55	76.41	108.46	121.39	139.70	281.05	0.216
B	10,000	128.77	26.99	72.66	110.17	122.82	140.59	249.58	0.210
C	10,000	98.51	22.45	56.73	82.05	93.57	109.48	225.89	0.228
Economic impact [k€]									
A	10,000	587.49	55.84	442.03	548.18	583.85	621.97	910.81	0.095
B	10,000	561.90	52.45	431.86	524.23	557.18	595.30	874.52	0.093
C	10,000	610.27	68.73	444.82	561.23	600.63	649.50	1,075.43	0.113
$\lambda = 1.00$									
Environmental impact [tCO_2eq]									
A	10,000	193.51	23.45	132.01	176.89	190.78	207.72	335.34	0.121
B	10,000	185.36	22.41	123.75	169.23	183.09	199.28	296.97	0.121
C	10,000	146.71	19.84	100.13	131.99	144.55	159.16	257.05	0.135
Economic impact [k€]									
A	10,000	594.08	54.66	446.75	555.22	590.38	628.50	914.68	0.092
B	10,000	563.11	53.08	430.94	525.12	558.45	597.29	874.21	0.094
C	10,000	612.20	69.47	444.82	562.00	603.31	652.44	1,085.83	0.113

Table B.2: Summary statistics for process-stage contributions grouped by impact type.

		Descriptive statistics						
Stage	n	μ	σ	min	$Q_{0.25}$	$Q_{0.50}$	$Q_{0.75}$	max
Environmental impact share [%]								
Forward stages								
New production	150,000	78.08	11.71	30.95	73.52	79.83	85.69	96.99
Refurbished production	150,000	10.96	8.19	0.00	5.82	9.88	14.54	54.70
Revitalization	150,000	7.47	5.41	0.34	4.25	5.98	8.23	31.98
Warehousing	150,000	0.02	0.02	0.00	0.00	0.02	0.03	0.29
Diesel transportation	150,000	0.25	0.37	0.00	0.00	0.09	0.23	1.60
Electric transportation	150,000	0.64	0.27	0.00	0.53	0.63	0.79	1.41
Reverse stages								
Collection	150,000	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Warehousing	150,000	0.00	0.00	0.00	0.00	0.00	0.00	0.02
Diesel transportation	150,000	0.06	0.10	0.00	0.00	0.01	0.05	0.50
Electric transportation	150,000	0.19	0.06	0.06	0.15	0.18	0.23	0.42
Disposal / end-of-life	150,000	2.44	1.21	0.59	1.68	2.19	2.96	13.18
Economic impact share [%]								
Forward stages								
New production	150,000	44.53	17.49	21.97	31.11	38.72	49.68	91.01
Refurbished production	150,000	37.53	17.47	0.00	32.92	42.80	50.69	66.23
Revitalization	150,000	5.48	3.11	0.66	3.37	4.62	6.49	20.01
Warehousing	150,000	0.14	0.21	0.00	0.00	0.02	0.21	2.41
Diesel transportation	150,000	0.06	0.50	0.00	0.00	0.00	0.00	5.30
Electric transportation	150,000	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Reverse stages								
Collection	150,000	3.39	1.20	1.51	2.45	2.83	4.68	6.50
Warehousing	150,000	0.01	0.01	0.00	0.00	0.00	0.01	0.39
Diesel transportation	150,000	2.05	3.49	0.00	0.00	0.22	1.35	13.14
Electric transportation	150,000	8.51	1.35	3.61	7.53	8.56	9.53	13.18
Disposal / end-of-life	150,000	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Table B.3: Signed step changes in expected quantity. Each entry reports Δx, the net change over the indicated λ step (in units). Positive values indicate net shifts toward the technology; negative values indicate shifts away.

		λ step			
Supplier	Technology	$0.00 \rightarrow 0.25$	$0.25 \rightarrow 0.50$	$0.50 \rightarrow 0.75$	$0.75 \rightarrow 1.00$
A	New production	+130.6	+10.0	+24.0	**+769.5**
A	Refurbishment	-149.8	-10.0	-23.9	**-753.6**
A	Revitalization	+19.2	-0.1	-0.0	-15.9
A	Warehousing	**-78.9**	-45.9	**-67.9**	-6.7
A	Diesel transport	+0.0	+0.0	+0.0	+2875.7[b]
A	Electric transport	-0.0	+0.0	+0.0	**-2875.7**[b]
A	Disposal	-19.2	+0.1	+0.0	+15.9
B	New production	+125.3	+62.6	+19.3	**+738.4**
B	Refurbishment	-141.4	-63.5	-19.0	**-726.7**
B	Revitalization	+16.1	+0.9	-0.3	-11.8
B	Warehousing	**-47.0**	-35.3	**-41.4**	-18.1
B	Diesel transport	+0.0	+0.0	+0.0	+3209.1[b]
B	Electric transport	+0.0	+0.0	+0.0	**-3209.1**[b]
B	Disposal	-16.1	-0.9	+0.3	+11.8
C	New production	+164.3	+2.2	+0.0	**+805.1**
C	Refurbishment	-183.0	-2.2	-0.0	**-791.6**
C	Revitalization	+18.7	+0.0	+0.0	-13.6
C	Warehousing[a]	-38.2	-0.7	-0.1	+742.0[c]
C	Diesel transport	+0.0	+0.0	+0.0	+3209.1[b]
C	Electric transport	+0.0	+0.0	+0.0	**-3209.1**[b]
C	Disposal	-18.7	-0.0	-0.0	+13.6

[a] Warehousing values for supplier C are affected by a run with zero cost (data-quality issue).
[b] Large transport shifts reflect a discrete switch between diesel and electric transportation in the highest λ step.
[c] Interpret this value with caution due to the issue in (a).

Appendix C

Auxiliary Contributions

C.1 ScenTrees: A Python Library for Scenario Trees

Multistage stochastic optimization requires discrete approximations of continuous stochastic processes in order to yield tractable deterministic equivalents. This appendix documents ScenTrees, a supportive Python library used in this thesis to construct scenario trees (for general path-dependent processes). The implementation follows the stochastic approximation framework of Pflug and Pichler (2015, 2016) and constitutes a Python port of the Julia package ScenTrees.jl (Kirui et al., 2020), with the repository archived in the bibliography (van Mastrigt, 2025, repository in bibliography). The aim of this section is not to reproduce the full theory, but to summarize the methodology at a level sufficient for interpretability and reproducibility.

C.1.1 Method summary

Let $\{\xi_t\}_{t=0}^{T}$ denote a stochastic process in $\mathbb{R}^d$, and let $\mathcal{T}$ be a scenario tree with a prescribed branching structure. ScenTrees seeks a discrete representation (node states and probabilities) that minimizes a multistage transportation distance (nested distance) between the empirical distribution of sampled trajectories and the discrete distribution induced by $\mathcal{T}$. The nested distance extends Wasserstein-type metrics to multistage settings by respecting the filtration (information) structure; in particular, deviations at early stages are propagated forward through the scenario structure (Pflug & Pichler, 2015, 2016).

Operationally, ScenTrees implements a stochastic approximation procedure akin to stochastic gradient descent in the space of discrete process measures. Given repeated samples of full trajectories, the tree (or lattice) is iteratively refined by matching samples to the currently closest discrete path (or stage-wise state) and updating the corresponding node values.

C.1.2 Stochastic approximation for trees

For a scenario tree, each leaf corresponds to a complete scenario path. In iteration k, a sample trajectory $\xi^{(k)} = (\xi_0^{(k)}, \dots, \xi_T^{(k)})$ is drawn from the data-generating mechanism (simulation model or data-driven sampler). The algorithm proceeds in four steps:

1. **Path matching:** identify the closest tree path (equivalently, the leaf whose root-to-leaf path minimizes the multistage distance to $\xi^{(k)}$).
2. **State update:** for each node on the matched path at stage t, update the node state in the direction of the sample realization at that stage.
3. **Probability update:** increase the visitation count of the matched leaf; leaf probabilities are obtained by normalization of visitation frequencies after K iterations.

4. **Exploration safeguard:** enforce a lower bound on branch probabilities early in the run to prevent premature loss of rarely visited branches.

The state update uses a generalized p-norm gradient with exponent parameter r (as in the Pflug–Pichler framework). Writing $\delta := x - \xi$ for the deviation between a current node state x and the matched sample state ξ (componentwise in $\mathbb{R}^d$), the update direction is computed as

$$g(\delta) = r \cdot \|\delta\|^{r-p} \cdot |\delta|^{p-1} \cdot \operatorname{sign}(\delta), \tag{C.1}$$

and the node state is updated by

$$x \leftarrow x - a_k \, g(\delta). \tag{C.2}$$

Following the stochastic approximation conditions of Robbins–Monro, ScenTrees employs an adaptive step size that decreases with repeated selection of the same leaf,

$$a_k = \frac{1}{30 + \rho_\ell}, \tag{C.3}$$

where ρ_ℓ denotes the visitation count (or associated mass) of the matched leaf ℓ. Intuitively, frequently selected paths are updated conservatively, while underrepresented paths receive comparatively larger corrections.

To reduce the risk of branch loss early in the optimization, the implementation includes a critical probability mechanism: leaf probabilities are prevented from falling below a time-varying threshold $p_{\text{crit}}(k)$, which decays with the iteration count. This maintains exploration in early iterations and stabilizes the resulting branching structure.

C.1.3 Lattices for Markovian processes

For Markovian processes, ScenTrees supports recombining scenario lattices. In contrast to trees, lattices store a set of states per stage and an explicit transition probability matrix between successive stages. The approximation proceeds stage-wise: at each stage t the sampled value $\xi_t^{(k)}$ is matched to the closest lattice state at that stage, and the corresponding state and transition counters are updated. Recombination reduces the number of nodes from exponential growth in T (trees) to roughly linear growth in the number of states per stage, making lattices preferable for long horizons when a Markov property is appropriate.

C.1.4 Implementation notes

The Python port mirrors the conceptual structure of ScenTrees.jl (Kirui et al., 2020), with explicit `Tree` and `Lattice` data structures storing node states and probability masses in NumPy arrays. The implementation separates (i) data structures, (ii) approximation routines (tree and lattice), and (iii) optional KDE-based trajectory generation.

C.1.5 Limitations

As with scenario generation methods in general, computational cost increases with the number of stages, leaves, and state dimension. Tree structures grow exponentially in the horizon length and therefore become impractical for large T unless branching is kept small; in such cases lattices provide a scalable alternative when a Markov assumption is acceptable. Moreover, approximation quality depends on user choices such as the branching structure, iteration budget, and the distance parameters (p, r); ScenTrees provides defaults but does not automate structure selection.

C.1.6 Repository

The complete source code for the Python implementation, including examples and documentation, is available in the archived repository cited in the bibliography (van Mastrigt, 2025).

C.2 Life-Cycle State-Space Formulation with Backlog

Up to this point, no formal state-space representation of the proposed formulation has been identified in the literature. Nevertheless, the model admits a natural interpretation within a discrete-time state-space framework. The main distinction from classical physical systems is that the state introduced below is primarily a *bookkeeping* state: it summarizes intertemporal effects of decisions rather than representing autonomous physical inertia.

State-space representations model dynamical systems using a state x_k, inputs u_k, disturbances w_k, and outputs y_k. For a linear time-invariant (LTI) or linear time-varying (LTV) system, a standard discrete-time form is

$$x_{k+1} = A_{ss,k}x_k + B_{ss,k}u_k + E_{ss,k}w_k, \tag{C.4}$$
$$y_k = C_{ss,k}x_k + D_{ss,k}u_k + F_{ss,k}w_k, \tag{C.5}$$

where the matrices may be time-varying in the LTV case. Such formulations expose temporal structure explicitly, support constraints naturally, and provide a basis for receding-horizon control and stochastic extensions.

Following Hülagü et al. (2025) in spirit (but using different notation), define the state x_k by

$$x_k := KA_{k-1}s_{k-1}, \qquad k \geq 2, \tag{C.6}$$

where A_k is the technology matrix at time k, s_k is the process scaling vector (decision variable), and K is a Leslie-type matrix encoding intertemporal propagation of flows. The initial condition is specified as

$$x_1 := 0. \tag{C.7}$$

With (C.6), the state update reads

$$x_{k+1} = KA_k s_k, \qquad k = 1, \ldots, T. \tag{C.8}$$

The variable x_k is interpreted as a *carry-in* contribution available in period k that is induced by decisions taken in period $k-1$ through (C.6). Thus, feasibility at time k depends on the current decision s_k, the current demand f_k, and the previously induced carry-in x_k. This timing convention aligns naturally with storage- or delay-type processes encoded in K.

Equation (C.8) implies that x_{k+1} does not explicitly depend on x_k (i.e., $A_{ss,k} = 0$). This is intentional: persistence is not introduced through autonomous state inertia but through explicit process decisions (e.g., storage-type processes) and, in the extension below, through a physical backlog state. The nilpotency of K further prevents unintended long-lived internal dynamics.

Define diagonal selection masks

$$E := \mathrm{diag}(\mathbf{e}), \qquad S := \mathrm{diag}(\mathbf{1} - \mathbf{e}), \qquad e \in \{0,1\}^N,$$

and define the residual

$$r_k := A_k s_k - f_k + x_k, \qquad k = 1, \ldots, T, \tag{C.9}$$

where f_k denotes demand. The baseline feasibility constraints are

$$E\, r_k = 0, \tag{C.10}$$
$$S\, r_k \geq 0, \tag{C.11}$$

for $k = 1, \ldots, T$. Componentwise, (C.10) enforces exact satisfaction when $e_i = 1$, while (C.11) enforces satisfaction with nonnegative slack when $e_i = 0$.

Impacts are modeled as

$$y_k = QB_k s_k. \tag{C.12}$$

Over a horizon T, define the cumulative criterion

$$J = \sum_{k=1}^{T} W\, Q B_k s_k. \tag{C.13}$$

For deterministic open-loop planning over a fixed horizon, the formulation can be written equivalently as a static multi-period program by eliminating x_k using (C.6). The state-space representation is retained because it (i) makes the temporal coupling explicit, (ii) supports receding-horizon re-optimization by conditioning on the current state, and (iii) provides a natural interface for stochastic formulations in which f_k (and potentially A_k or capacities) are uncertain and decisions are adapted over time.

C.2.1 Backlog extension

Requiring all demands to be met within each period may be overly restrictive in practice: capacity limitations, lead-times, or disruptions may prevent on-time fulfillment. A standard remedy is a *backlog* (backorder) state that carries unmet demand forward in time, thereby preserving feasibility and enabling explicit trade-offs between service and impact/cost.

Let $b_k \in \mathbb{R}^N_{\geq 0}$ denote the backlog at time k. Not all rows necessarily permit backlog; introduce a binary selector $\mathbf{e}_b \in \{0,1\}^N$ and the diagonal mask

$$U := \operatorname{diag}(\mathbf{e}_b). \tag{C.14}$$

To maintain consistency with the baseline masks, backlog is permitted only on slack-permitted rows:

$$\mathbf{e}_b \leq \mathbf{1} - \mathbf{e} \quad \text{(componentwise)} \qquad \Longleftrightarrow \qquad U \preceq S. \tag{C.15}$$

For $k = 1, \ldots, T$,

$$s_k \in \mathbb{R}_+^{P_{\text{con}}} \times \mathbb{Z}_+^{P_{\text{dis}}}, \qquad x_1 := 0, \qquad b_1 \text{ given (often 0)}.$$

The state trajectories are treated as decision variables subject to the system dynamics:

$$x_k \in \mathbb{R}^N_+, \quad k = 1, \ldots, T+1, \qquad b_k \in \mathbb{R}^N_+, \quad k = 1, \ldots, T+1.$$

On permitted rows, backlog accumulates unmet demand:

$$b_{k+1} = b_k + U f_k - U(A_k s_k + x_k), \qquad k = 1, \ldots, T. \tag{C.16}$$

Equivalently, using (C.9), (C.16) can be written as

$$b_{k+1} = b_k - U r_k, \qquad k = 1, \ldots, T. \tag{C.17}$$

Define the augmented state $z_k := \begin{bmatrix} x_k \\ b_k \end{bmatrix}$. Combining (C.8) and (C.16) yields the LTV update

$$\begin{bmatrix} x_{k+1} \\ b_{k+1} \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ -U & I \end{bmatrix} \begin{bmatrix} x_k \\ b_k \end{bmatrix} + \begin{bmatrix} KA_k \\ -UA_k \end{bmatrix} s_k + \begin{bmatrix} 0 \\ U \end{bmatrix} f_k, \qquad k = 1, \ldots, T, \tag{C.18}$$

where demand f_k acts as an exogenous input/disturbance.

Allow shortfalls only on backlog-permitted rows by modifying the baseline slack constraint:

$$E\, r_k = 0, \qquad k = 1, \ldots, T, \tag{C.19}$$

$$(S - U)\, r_k \geq 0, \qquad k = 1, \ldots, T, \tag{C.20}$$

and enforce backlog support only where allowed:

$$(I - U)\, b_k = 0, \qquad k = 1, \ldots, T+1. \tag{C.21}$$

Capacity constraints are imposed as in the baseline model, e.g.

$$A_k^{\text{out}}\, s_k \le c_k, \qquad k = 1, \ldots, T, \tag{C.22}$$

where c_k denotes available capacity.

To penalize backlog analogously to process activity, introduce an auxiliary variable $\ell_k \in \mathbb{R}_+^N$ and define

$$\tilde{s}_k := \begin{bmatrix} s_k \\ \ell_k \end{bmatrix}, \qquad \tilde{B}_k := \begin{bmatrix} B_k & B_k^{\text{bl}} \end{bmatrix}.$$

A natural choice is to penalize backlog carried into the next period:

$$\ell_k = b_{k+1}, \qquad k = 1, \ldots, T, \tag{C.23}$$

which induces an implicit terminal penalty on b_{T+1}. (Alternatively one may set $\ell_k = b_k$ to penalize backlog present during period k.)

To keep the material balance unaffected by ℓ_k, define

$$\tilde{A}_k := \begin{bmatrix} A_k & 0_{N\times N} \end{bmatrix}, \qquad \tilde{B}_k := \begin{bmatrix} B_k & B_k^{\text{bl}} \end{bmatrix}.$$

Then (C.18) extends to

$$\begin{bmatrix} x_{k+1} \\ b_{k+1} \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ -U & I \end{bmatrix} \begin{bmatrix} x_k \\ b_k \end{bmatrix} + \begin{bmatrix} K\tilde{A}_k \\ -U\tilde{A}_k \end{bmatrix} \tilde{s}_k + \begin{bmatrix} 0 \\ U \end{bmatrix} f_k, \qquad k = 1, \ldots, T. \tag{C.24}$$

The corresponding output becomes

$$y_k = Q\, \tilde{B}_k\, \tilde{s}_k. \tag{C.25}$$

Two common terminal treatments are: (i) *soft terminal service*, achieved by penalizing b_{T+1} through (C.23), and/or (ii) *hard terminal service*, enforced via

$$b_{T+1} = 0, \tag{C.26}$$

(or row-wise $U b_{T+1} = 0$) when full fulfillment by the end of the horizon is required. The choice depends on whether the horizon is interpreted as an operational planning window (soft terminal) or a complete project lifecycle (hard terminal).

The backlog state prevents infeasibility under temporary shortfalls and provides a natural service-level measure. Under uncertainty in f_k (and potentially A_k or c_k), the augmented state z_k supports multistage (scenario-tree) formulations and receding-horizon policies $s_k = \pi_k(z_k)$, since feasibility and state evolution depend on current state, current decision, and the realized disturbance.

Finally, an optimization problem over the horizon can be posed as

$$\min_{\{\tilde{s}_k\}_{k=1}^T} \sum_{k=1}^{T} W\, Q\, \tilde{B}_k\, \tilde{s}_k, \tag{C.27}$$

subject to the dynamics (C.24), feasibility constraints (C.19)–(C.21), capacity constraints (C.22), terminal conditions (soft and/or hard) (C.23)–(C.26), and the variable domain constraints.

Bibliography

Apap, R. M., & Grossmann, I. (2017). Models and computational strategies for multistage stochastic programming under endogenous and exogenous uncertainties. *Comput. Chem. Eng.*, *103*, 233–274. https://doi.org/10.1016/j.compchemeng.2016.11.011

Babaï, M., Syntetos, A., & Teunter, R. (2014). Intermittent demand forecasting: An empirical study on accuracy and the risk of obsolescence. *International Journal of Production Economics*, *157*, 212–219. https://doi.org/10.1016/j.ijpe.2014.08.019

Bairamzadeh, S., Pishvaee, M. S., & Saidi-Mehrabad, M. (2015). Multiobjective robust possibilistic programming approach to sustainable bioethanol supply chain design under multiple uncertainties. *Industrial & Engineering Chemistry Research*, *55*(1), 237–256. https://doi.org/10.1021/acs.iecr.5b02875

Bamber, N., Turner, I., Arulnathan, V., Li, Y., Zargar, S., Smart, A., & Pelletier, N. (2019). Comparing sources and analysis of uncertainty in consequential and attributional life cycle assessment: Review of current practice and recommendations. *The International Journal of Life Cycle Assessment*, *25*(1), 168–180. https://doi.org/10.1007/s11367-019-01663-1

Bandara, K., Bergmeir, C., & Smyl, S. (2020). Forecasting across time series databases using recurrent neural networks on groups of similar series: A clustering approach. *Expert Systems With Applications*, *140*, 112896. https://doi.org/10.1016/j.eswa.2019.112896

Barahmand, Z., & Eikeland, M. S. (2022). Life cycle assessment under uncertainty: a scoping review. *World*, *3*(3), 692–717. https://doi.org/10.3390/world3030039

Barbosa-Póvoa, A., Silva, C., & Carvalho, A. (2017). Opportunities and challenges in sustainable supply chain: An operations research perspective. *Eur. J. Oper. Res.*, *268*, 399–431. https://doi.org/10.1016/j.ejor.2017.10.036

Beltrán, A. M., Heijungs, R., Guinée, J. B., & Tukker, A. (2015). A pseudo-statistical approach to treat choice uncertainty: the example of partitioning allocation methods. *The International Journal of Life Cycle Assessment*, *21*(2), 252–264. https://doi.org/10.1007/s11367-015-0994-4

Birge, J. R., & Louveaux, F. (2011). *Introduction to stochastic programming*. Springer New York. https://doi.org/10.1007/978-1-4614-0237-4

Boland, N., Christiansen, J., Dandurand, B., Eberhard, A. C., Linderoth, J. T., Luedtke, J. R., & Oliveira, F. (2017). Combining progressive hedging with a frank-wolfe method to compute lagrangian dual bounds in stochastic mixed-integer programming. *SIAM J. Optim.*, *28*, 1312–1336. https://doi.org/10.1137/16m1076290

Brandenburg, M., Govindan, K., Sarkis, J., & Seuring, S. (2014). Quantitative models for sustainable supply chain management: Developments and directions. *Eur. J. Oper. Res.*, *233*, 299–312. https://doi.org/10.1016/j.ejor.2013.09.032

Budzinski, M., Sisca, M., & Thrän, D. (2019). Consequential lca and lcc using linear programming: An illustrative example of biorefineries. *The International Journal of Life Cycle Assessment*, *24*(12), 2191–2205. https://doi.org/10.1007/s11367-019-01650-6

Bynum, M. L., Hackebeil, G. A., Hart, W. E., Laird, C. D., Nicholson, B. L., Siirola, J. D., Watson, J.-P., & Woodruff, D. L. (2021). *Pyomo–optimization modeling in python* (Third, Vol. 67). Springer Science & Business Media.

Calvin, K., Dasgupta, D., Krinner, G., Mukherji, A., Thorne, P. W., Trisos, C., Romero, J., Aldunce, P., Barrett, K., Blanco, G., Cheung, W. W., Connors, S., Denton, F., Diongue-Niang, A., Dodman, D., Garschagen, M., Geden, O., Hayward, B., Jones, C., … Ha, M. (2023, July). *IPCC, 2023: Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change [Core Writing Team, H. Lee and J. Romero (eds.)]. IPCC, Geneva, Switzerland.* (P. Arias, M. Bustamante, I. Elgizouli, G. Flato, M. Howden, C. Méndez-Vallejo, J. J. Pereira, R. Pichs-Madruga, S. K. Rose, Y. Saheb, R. Sánchez Rodríguez, D. Ürge-Vorsatz, C. Xiao, N. Yassaa, J. Romero, J. Kim, E. F. Haites, Y. Jung, R. Stavins, … C. Péan, Eds.). Intergovernmental Panel on Climate Change (IPCC). https://doi.org/10.59327/ipcc/ar6-9789291691647

Carøe, C. C., & Tind, J. (1998). L-shaped decomposition of two-stage stochastic programs with integer recourse. *Mathematical Programming*, *83*(1–3), 451–464. https://doi.org/10.1007/bf02680570

Carrillo Pineda, A., Watson, E., Castro, E., White, E., Rangel Villasana, F., Borjigin-Wang, E., Frédeau, M., Nielsen, J., Ponce de León Baridó, P., Von Koeller, E., Collins, N., Verde, M., & Lawrence, B. (2023, February). *Catalyzing Value Chain Decarbonization: Corporate Survey Results* (tech. rep.) (Report based on stakeholder engagement survey conducted by SBTi with support from BCG). Science Based Targets initiative (SBTi) and Boston Consulting Group (BCG).

Ciroth, A., Müller, S., Weidema, B., & Lesage, P. (2013). Empirically based uncertainty factors for the pedigree matrix in ecoinvent. *The International Journal of Life Cycle Assessment*, *21*(9), 1338–1348. https://doi.org/10.1007/s11367-013-0670-5

Cucurachi, S., Borgonovo, E., & Heijungs, R. (2015). A protocol for the global sensitivity analysis of impact assessment models in life cycle assessment. *Risk Analysis*, *36*(2), 357–377. https://doi.org/10.1111/risa.12443

Delft High Performance Computing Centre. (2024). *Delftblue supercomputer (phase 2)* [ARK: ark:/44463/DelftBluePhase2]. Retrieved December 30, 2025, from https://www.tudelft.nl/dhpc/ark:/44463/DelftBluePhase2

Ding, T., Hu, Y., & Bie, Z. (2018). Multi-stage stochastic programming with nonanticipativity constraints for expansion of combined power and natural gas systems. *IEEE Transactions on Power Systems*, *33*, 317–328. https://doi.org/10.1109/tpwrs.2017.2701881

Dorland, J., & Jørgensen, M. S. (2025). Unpacking repair: A comparative study of socio-technical factors influencing consumer decisions in the circular economy across three product categories. *Circular Economy and Sustainability*, *5*(5), 3991–4026. https://doi.org/10.1007/s43615-025-00596-2

Dupacová, J., Consigli, G., & Wallace, S. (2000). Scenarios for multistage stochastic programs. *Annals of Operations Research*, *100*, 25–53. https://doi.org/10.1023/a:1019206915174

Dupačová, J., Gröwe-Kuska, N., & Römisch, W. (2003). Scenario reduction in stochastic programming. *Mathematical Programming*, *95*(3), 493–511. https://doi.org/10.1007/s10107-002-0331-0

Dyer, M., & Stougie, L. (2006). Computational complexity of stochastic programming problems. *Mathematical Programming*, *106*, 423–432. https://doi.org/10.1007/s10107-005-0597-0

Eurostat. (2022a). Goods transported by group of goods - from 2008 onwards based on nst 2007. https://doi.org/10.2908/rail_go_grpgood

Eurostat. (2022b). Transport by type of good (from 2007 onwards with nst2007). https://doi.org/10.2908/iww_go_atygo

Eurostat. (2023). Road freight transport by type of goods and type of transport (t, tkm) - annual data. https://doi.org/10.2908/road_go_ta_tg

Fildes, R., Goodwin, P., Lawrence, M., & Nikolopoulos, K. (2009). Effective forecasting and judgmental adjustments: An empirical evaluation and strategies for improve-

ment in supply-chain planning. *International Journal of Forecasting*, *25*(1), 3–23. https://doi.org/10.1016/j.ijforecast.2008.11.010

Frolov, M., Tanchenko, S., & Ohluzdina, L. (2024). Parameter estimation of the weibull distribution in modeling the reliability of technical objects. *Journal of Engineering Sciences*. https://doi.org/10.21272/jes.2024.11(1).a1

Gao, J., & You, F. (2017). Modeling framework and computational algorithm for hedging against uncertainty in sustainable supply chain design using functional-unit-based life cycle optimization. *Comput. Chem. Eng.*, *107*, 221–236. https://doi.org/10.1016/j.compchemeng.2017.05.021

Gómez, Y., Gallardo, D., Marchant, C., Sánchez, L., & Bourguignon, M. (2023). An in-depth review of the weibull model with a focus on various parameterizations. *Mathematics*. https://doi.org/10.3390/math12010056

Gooijer, J. G. D., & Hyndman, R. J. (2006). 25 years of time series forecasting. *International Journal of Forecasting*, *22*(3), 443–473. https://doi.org/10.1016/j.ijforecast.2006.01.001

Groen, E., Bokkers, E., Heijungs, R., & Boer, I. d. (2016). Methods for global sensitivity analysis in life cycle assessment. *The International Journal of Life Cycle Assessment*, *22*(7), 1125–1137. https://doi.org/10.1007/s11367-016-1217-3

Guillén-Gosálbez, G., & Grossmann, I. E. (2008). Optimal design and planning of sustainable chemical supply chains under uncertainty. *AIChE Journal*, *55*(1), 99–121. https://doi.org/10.1002/aic.11662

Gurobi Optimization, LLC. (2025). Gurobi Optimizer Reference Manual. https://www.gurobi.com/

Hart, W. E., Watson, J.-P., & Woodruff, D. L. (2011). Pyomo: Modeling and solving mathematical programs in python. *Mathematical Programming Computation*, *3*(3), 219–260.

Heijungs, R., Settanni, E., & Guinée, J. (2013). Toward a computational structure for life cycle sustainability analysis: Unifying lca and lcc. *The International Journal of Life Cycle Assessment*, *18*, 1722–1733. https://doi.org/10.1007/s11367-012-0461-4

Heijungs, R. (2019). On the number of monte carlo runs in comparative probabilistic lca. *The International Journal of Life Cycle Assessment*, *25*(2), 394–402. https://doi.org/10.1007/s11367-019-01698-4

Heijungs, R. (2023). Analysis and remediation of the confusing specification of the log-normal distribution. *The International Journal of Life Cycle Assessment*, *29*(3), 537–554. https://doi.org/10.1007/s11367-023-02249-8

Heijungs, R. (2024a). *Probability, statistics and life cycle assessment: Guidance for dealing with uncertainty and sensitivity*. Springer International Publishing. https://doi.org/10.1007/978-3-031-49317-1

Heijungs, R. (2024b). Uncertainty and sensitivity analysis in life cycle assessment. In *Encyclopedia of sustainable technologies* (pp. 235–248). Elsevier. https://doi.org/10.1016/b978-0-323-90386-8.00039-5

Heijungs, R., & Huijbregts, M. A. J. (2004). A review of approaches to treat uncertainty in lca. *Proceedings of the 2nd International Congress on Environmental Modelling and Software*.

Heijungs, R., & Suh, S. (2002). *The computational structure of life cycle assessment*. Kluwer Academic Publishers.

Heitsch, H., & Römisch, W. (2009). Scenario tree modeling for multistage stochastic programs. *Mathematical Programming*, *118*, 371–406. https://doi.org/10.1007/s10107-007-0197-2

Heller-Ono, A. (2020). Predictive analytics for chair asset management. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *64*(1), 1223–1227. https://doi.org/10.1177/1071181320641291

Herth, A., & Blok, K. (2022). Quantifying universities' direct and indirect carbon emissions – the case of Delft University of Technology. *International Journal of Sustainability in Higher Education*, *24*(9), 21–52.

Huijbregts, M. A. J. (1998). Application of uncertainty and variability in lca. *The International Journal of Life Cycle Assessment*, *3*(5). https://doi.org/10.1007/bf02979835

Huijbregts, M. A. J., Gilijamse, W., Ragas, A. M., & Reijnders, L. (2003). Evaluating uncertainty in environmental life-cycle assessment. a case study comparing two insulation options for a dutch one-family dwelling. *Environmental Science & Technology*, *37*(11), 2600–2608. https://doi.org/10.1021/es020971+

Hülagü, S., Dullaert, W., Eruguz, A. S., Heijungs, R., & Inghels, D. (2025). Integrating life cycle assessment into supply chain optimization (A. Barbaresi, Ed.). *PLOS ONE*, *20*(1), e0316710. https://doi.org/10.1371/journal.pone.0316710

Hung, M., & Ma, H. (2008). Quantifying system uncertainty of life cycle assessment based on monte carlo simulation. *The International Journal of Life Cycle Assessment*, *14*(1), 19–27. https://doi.org/10.1007/s11367-008-0034-8

Igos, E., Benetto, E., Meyer, R., Baustert, P., & Othoniel, B. (2018). How to treat uncertainties in life cycle assessment studies? *The International Journal of Life Cycle Assessment*, *24*(4), 794–807. https://doi.org/10.1007/s11367-018-1477-1

Imbeault-Tétreault, H., Jolliet, O., Deschênes, L., & Rosenbaum, R. K. (2013). Analytical propagation of uncertainty in life cycle assessment using matrix formulation. *Journal of Industrial Ecology*, *17*(4), 485–492. https://doi.org/10.1111/jiec.12001

Iooss, B., & Lemaître, P. (2015). A review on global sensitivity analysis methods. *Operations Research/Computer Science Interfaces Series*, 101–122. https://doi.org/10.1007/978-1-4899-7547-8_5

ISO 14040:2006(E). (2006). *Environmental management – Life cycle assessment – Principles and framework* (Standard). International Organization for Standardization. Geneva, CH.

Jayarathna, C., Agdas, D., Dawes, L., & Yigitcanlar, T. (2021). Multi-objective optimization for sustainable supply chain and logistics: A review. *Sustainability*. https://doi.org/10.3390/su132413617

Jetti, R., & Dhar, D. (2024). Product durability: A systematic literature review. *Archives of Design Research*, *37*(5), 93–128. https://doi.org/10.15187/adr.2024.11.37.5.93

Jiang, R. (2020). A novel two-fold sectional approximation of renewal function and its applications. *Reliab. Eng. Syst. Saf.*, *193*, 106624. https://doi.org/10.1016/j.ress.2019.106624

Jiang, R., & Murthy, D. (2011). A study of weibull shape parameter: Properties and significance. *Reliab. Eng. Syst. Saf.*, *96*, 1619–1626. https://doi.org/10.1016/j.ress.2011.09.003

Joe, H. (2014, June). *Dependence modeling with copulas*. Chapman; Hall/CRC. https://doi.org/10.1201/b17116

Kaynak, E., Piri, I., & Das, O. (2025). Revisiting the basics of life cycle assessment and lifecycle thinking. *Sustainability*. https://doi.org/10.3390/su17167444

Kessy, A., Lewin, A., & Strimmer, K. (2015). Optimal whitening and decorrelation. *The American Statistician*, *72*, 309–314. https://doi.org/10.1080/00031305.2016.1277159

Kim, A., Mutel, C., & Hellweg, S. (2025). Global sensitivity analysis of correlated uncertainties in life cycle assessment. *Journal of Industrial Ecology*, *29*(4), 1090–1104. https://doi.org/10.1111/jiec.70036

Kirui, K., Pichler, A., & Pflug, G. C. (2020). Scentrees.jl: A julia package for generating scenario trees and scenario lattices for multistage stochastic programming. *Journal of Open Source Software*, *5*(46), 1912. https://doi.org/10.21105/joss.01912

Kleywegt, A., Shapiro, A., & Homem-de-Mello, T. (2002). The sample average approximation method for stochastic discrete optimization. *SIAM J. Optim.*, *12*, 479–502. https://doi.org/10.1137/s1052623499363220

Knueven, B., Mildebrath, D., Muir, C., Siirola, J. D., Watson, J.-P., & Woodruff, D. L. (2023). A parallel hub-and-spoke system for large-scale scenario-based optimization under uncertainty. *Math. Prog. Comp.*, *15*, 591–619.

Kokhuis, K., Van den Berghe, K., & Straub, A. (2025). Implementing circular practices in a construction clients' organisation. strategic interventions on intra-organisational barriers for operationalising tu delft's circular campus ambitions. *Journal of Corporate Real Estate*, *27*(4), 281–301. https://doi.org/10.1108/jcre-11-2024-0038

Kumar, K., & Kumar, A. (2024). Application of optimization models in sustainable supply chain management: A systematic review based on prisma guidelines. *Process Integration and Optimization for Sustainability*, *9*, 763–790. https://doi.org/10.1007/s41660-024-00455-x

Leal Filho, W., Shiel, C., do Paço, A., & Brandli, L. (2015). Putting sustainable development in practice: Campus greening as a tool for institutional sustainability efforts. In *Sustainability in higher education* (pp. 1–19). Elsevier.

Lee, H. L., Padmanabhan, V., & Whang, S. (1997). The bullwhip effect in supply chains. *Sloan Management Review*, *38*(3), 93–102.

Lozano-Oviedo, J., Cortés, C. E., & Rey, P. A. (2024). Sustainable closed-loop supply chains and their optimization models: A review of the literature. *Clean Technologies and Environmental Policy*, *26*, 999–1023. https://doi.org/10.1007/s10098-023-02730-w

Mak, W., Morton, D. P., & Wood, R. K. (1999). Monte carlo bounding techniques for determining solution quality in stochastic programs. *Operations Research Letters*, *24*(1-2), 47–56. https://doi.org/10.1016/s0167-6377(98)00054-6

McCool, J. (2012). Using the weibull distribution: Reliability, modeling, and inference. https://doi.org/10.1002/9781118351994

Moreno Ruiz, E., Valsasina, L., Brunner, F., Symeonidis, A., FitzGerald, D., Treyer, K., Bourgault, G., & Wernet, G. (2018). *Documentation of changes implemented in the ecoinvent database v3.5* (Technical report). ecoinvent Association. Zurich, Switzerland. Retrieved June 1, 2023, from https://www.ecoinvent.org/files/change_report_v3_5_20180823.pdf

Mota, B., Gomes, M. I., Carvalho, A., & Barbosa-Póvoa, A. (2017). Sustainable supply chains: An integrated modeling approach under uncertainty. *Omega*. https://doi.org/10.1016/j.omega.2017.05.006

Müller, S., Lesage, P., Ciroth, A., Mutel, C., Weidema, B. P., & Samson, R. (2014). The application of the pedigree approach to the distributions foreseen in ecoinvent v3. *The International Journal of Life Cycle Assessment*, *21*(9), 1327–1337. https://doi.org/10.1007/s11367-014-0759-5

Mutel, C. (2017). Brightway: An open source framework for life cycle assessment. *Journal of Open Source Software*, *2*(12), 236. https://doi.org/10.21105/joss.00236

Mutel, C. (2025). Stats-arrays [Version 1.0]. https://pypi.org/project/stats-arrays/

Nemhauser, G., & Wolsey, L. (1988, June). *Integer and combinatorial optimization*. Wiley. https://doi.org/10.1002/9781118627372

Paris Agreement, 2015, https://treaties.un.org/pages/ViewDetails.aspx?src=TREATY%5C&mtdsg_no=XXVII-7-d%5C&chapter=27%5C&clang=_en

Park, H., & Hanasusanto, G. A. (2024). Sample complexity of data-driven multistage stochastic programming under markovian uncertainty. https://arxiv.org/abs/2412.19299

Peyré, G., & Cuturi, M. (2020). Computational optimal transport. https://arxiv.org/abs/1803.00567

Pflug, G. C., & Pichler, A. (2015). Dynamic generation of scenario trees. *Computational Optimization and Applications*, *62*(3), 641–668. https://doi.org/10.1007/s10589-015-9758-0

Pflug, G. C., & Pichler, A. (2016). From empirical observations to tree models for stochastic optimization: Convergence properties. *SIAM Journal on Optimization*, *26*(3), 1715–1740. https://doi.org/10.1137/15m1043376

Piepenbrink, H., Flämig, H., & Menger, A. (2025). CO2e life-cycle assessment: Twin comparison of battery–electric and diesel heavy-duty tractor units with real-world data. *Future Transportation*, *5*(1), 12. https://doi.org/10.3390/futuretransp5010012

Qin, Y., & Suh, S. (2016). What distribution function do life cycle inventories follow? *The International Journal of Life Cycle Assessment*, *22*(7), 1138–1145. https://doi.org/10.1007/s11367-016-1224-4

Rinne, H. (2008, November). *The weibull distribution*. Chapman; Hall/CRC. https://doi.org/10.1201/9781420087444

Ross, S. A., & Cheah, L. (2018). Uncertainty quantification in life cycle assessments: exploring distribution choice and greater data granularity to characterize product use. *Journal of Industrial Ecology*, *23*(2), 335–346. https://doi.org/10.1111/jiec.12742

Rožanec, J. M., & Mladenić, D. (2021). Reframing demand forecasting: A two-fold approach for lumpy and intermittent demand. *arXiv*. https://doi.org/10.3390/su14159295

Ruiz-Femenia, R., Guillén-Gosálbez, G., Jiménez, L., & Caballero, J. A. (2013). Multi-objective optimization of environmentally conscious chemical supply chains under demand uncertainty. *Chemical Engineering Science*, *95*, 1–11. https://doi.org/10.1016/j.ces.2013.02.054

Seyedan, M., & Mafakheri, F. (2020). Predictive big data analytics for supply chain demand forecasting: Methods, applications, and research opportunities. *Journal of Big Data*, *7*(1). https://doi.org/10.1186/s40537-020-00329-2

Shapiro, A. (2011). Analysis of stochastic dual dynamic programming method. *Eur. J. Oper. Res.*, *209*, 63–72. https://doi.org/10.1016/j.ejor.2010.08.007

Shapiro, A. (2007). Stochastic programming approach to optimization under uncertainty. *Mathematical Programming*, *112*, 183–220. https://doi.org/10.1007/s10107-006-0090-4

Simchi-Levi, D., Kaminsky, P., & Simchi-Levi, E. (2022). *Designing and managing the supply chain: Concepts, strategies and case studies* (Fourth edition) [ISBN-10: 1259997812]. McGraw-Hill LLC.

Sonderegger, T. (2024, November 19). *Implementation of life cycle impact assessment methods in the ecoinvent database v3.11* (Technical Report). ecoinvent Association. Zürich, Switzerland.

Suh, S., Weidema, B., Schmidt, J., & Heijungs, R. (2010). Generalized make and use framework for allocation in life cycle assessment. *Journal of Industrial Ecology*, *14*. https://doi.org/10.1111/j.1530-9290.2010.00235.x

Syntetos, A., Babai, M., & Gardner, E. S. (2015). Forecasting intermittent inventory demands: Simple parametric methods vs. bootstrapping. *Journal of Business Research*, *68*, 1746–1752. https://doi.org/10.1016/j.jbusres.2015.03.034

Syntetos, A., & Boylan, J. (2005). The accuracy of intermittent demand estimates. *International Journal of Forecasting*, *21*, 303–314. https://doi.org/10.1016/j.ijforecast.2004.10.001

Syntetos, A. A., Boylan, J., & Croston, J. D. (2005). On the categorization of demand patterns. *Journal of the Operational Research Society*, *56*, 495–503. https://doi.org/10.1057/palgrave.jors.2601841

Teunter, R., & Duncan, L. (2009). Forecasting intermittent demand: A comparative study. *Journal of the Operational Research Society*, *60*, 321–329. https://doi.org/10.1057/palgrave.jors.2602569

Torres Figueroa, J. J., Li, C., Apap, R. M., & Grossmann, I. (2022). A review on the performance of linear and mixed integer two-stage stochastic programming software. *Algorithms*, *15*, 103. https://doi.org/10.3390/a15040103

Türkmen, A. C., Januschowski, T., Wang, Y., & Cemgil, A. T. (2021). Forecasting intermittent and sparse time series: A unified probabilistic framework via deep renewal processes. *PLoS ONE*, *16*. https://doi.org/10.1371/journal.pone.0259764

van Gameren, D., & van den Dobbelsteen, A. (2023). The sustainable campus: Working towards a carbon-neutral university. In S. Roaf & W. Finlayson (Eds.), *Measuring net zero: Carbon accounting for buildings and communities* (pp. 273–285). Ecohouse Initiative Ltd.

van Mastrigt, P. G. (2025). *Scentrees-python: A python port of scentrees.jl for scenario tree and lattice approximation* [Python reimplementation and extension of the ScenTrees.jl package]. https://github.com/philmas/scentrees-python

van den Dobbelsteen, A., van Gameren, D., Roeling, M., & van Mourik, T. (2023). Practising what you teach and preach: Working for a sustainable university is a verb. In M. Ban, N. Duić, D. R. Schneider, D. Astiaso Garcia, G. Barone, S. Boldyryev, A. Buonomano, F. Calise, & P. Colbertaldo (Eds.), *Proceedings of the 18th conference on sustainable development of energy, water and environment systems (sdewes 2023)* (p. 85). University of Zagreb. https://www.dubrovnik2023.sdewes.org/

van den Dobbelsteen, A., & van Gameren, D. (2022). *Sustainable TU Delft - Vision, Ambition and Action Plan for a Climate University*. Delft University of Technology.

Villani, C. (2009). The wasserstein distances. In *Optimal transport: Old and new* (pp. 93–111). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-540-71050-9_6

Volvo Trucks. (2022, January). Mise à l'essai du volvo electric : D'excellents résultats en termes d'autonomie et de rendement énergétique [Press release reporting a measured consumption of 1.1 kWh/km for a Volvo FH Electric at 40 t gross weight in independent testing]. https://www.volvotrucks.fr/fr-fr/news/press-releases/2022/jan/volvos-heavy-duty-electric-truck-is-put-to-the-test-excels-in-both-range-and-energy-efficiency.html

Watson, J., Woodruff, D. L., & Hart, W. (2012). Pysp: Modeling and solving stochastic programs in python. *Mathematical Programming Computation*, *4*, 109–149. https://doi.org/10.1007/s12532-012-0036-1

Weglarczyk, S. (2018). Kernel density estimation and its application. *ITM Web of Conferences*, *23*, 00037. https://doi.org/10.1051/itmconf/20182300037

Wei, W., Larrey-Lassalle, P., Faure, T., Dumoulin, N., Roux, P., & Mathias, J. (2014). How to conduct a proper sensitivity analysis in life cycle assessment: taking into account correlations within lci data and interactions within the lca calculation model. *Environmental Science & Technology*, *49*(1), 377–385. https://doi.org/10.1021/es502128k

Weibull, W. (1951). A statistical distribution function of wide applicability. *Journal of Applied Mechanics*, *18*(3), 293–297. https://doi.org/10.1115/1.4010337

Weidema, B. P., & Wesnæs, M. S. (1996). Data quality management for life cycle inventories—an example of using data quality indicators. *Journal of Cleaner Production*, *4*(3–4), 167–174. https://doi.org/10.1016/s0959-6526(96)00043-1

Wernet, G., Bauer, C., Steubing, B., Reinhard, J., Moreno-Ruiz, E., & Weidema, B. (2016). The ecoinvent database version 3 (part i): Overview and methodology. *The International Journal of Life Cycle Assessment*, *21*(9), 1218–1230. https://doi.org/10.1007/s11367-016-1087-8

Willemain, T., Smart, C. N., & Schwarz, H. (2004). A new approach to forecasting intermittent demand for service parts inventories. *International Journal of Forecasting*, *20*, 375–387. https://doi.org/10.1016/s0169-2070(03)00013-x

Williams, E., Weber, C. L., & Hawkins, T. R. (2009). Hybrid framework for managing uncertainty in life cycle inventories. *Journal of Industrial Ecology*, *13*(6), 928–944. https://doi.org/10.1111/j.1530-9290.2009.00170.x

You, F., Tao, L., Graziano, D., & Snyder, S. W. (2011). Optimal design of sustainable cellulosic biofuel supply chains: multiobjective optimization coupled with life cycle assessment and input–output analysis. *AIChE Journal*, *58*(4), 1157–1180. https://doi.org/10.1002/aic.12637

Yue, D., Kim, M. A., & You, F. (2013). Design of sustainable product systems and supply chains with life cycle optimization based on functional unit: General modeling framework, mixed-integer nonlinear programming algorithms and case study on hydrocarbon biofuels. *ACS Sustainable Chemistry & Engineering*, *1*, 1003–1014. https://doi.org/10.1021/sc400080x

Zelm, R. v., & Huijbregts, M. A. J. (2013). Quantifying the trade-off between parameter and model structure uncertainty in life cycle impact assessment. *Environmental Science & Technology*, *47*(16), 9274–9280. https://doi.org/10.1021/es305107s