

UNIVERSITÀ DEGLI STUDI DI GENOVA
SCUOLA POLITECNICA

DIME

**Dipartimento di Ingegneria Meccanica, Energetica,
Gestionale e dei Trasporti**



TESI DI LAUREA MAGISTRALE
IN
INGEGNERIA GESTIONALE

**Circular Business Models for Critical Raw Material
Recovery: A Scenario-Based DSS with an Application
in the Offshore Wind Sector**

Relatore:

Prof. Stefania Testa

Correlatore:

Prof. Kepa Mendibil

Allievo:

Matilde Delfino

Ottobre 2026

Modelli di Business Circolari per il Recupero delle Materie Prime Critiche: un DSS Basato su Scenari con un'Applicazione nel Settore Eolico Offshore

Sommario

La transizione verso le energie rinnovabili ha aumentato significativamente la domanda di Materie Prime Critiche (CRMs), esponendo i produttori a rischi di approvvigionamento e volatilità dei prezzi. I Circular Business Models (CBMs) offrono modalità alternative di organizzazione della proprietà, delle responsabilità e del controllo a fine vita di un asset, ma gli archetipi orientati al prodotto, all'uso e al risultato comportano trade-off differenti. La letteratura offre un supporto strutturato limitato per confrontare tali alternative dalla prospettiva di un OEM in presenza di condizioni di mercato dei CRM variabili.

La tesi sviluppa un Decision Support System (DSS) basato su scenari per confrontare i tre archetipi. Seguendo un approccio di Design Science Research, le variabili decisionali derivate dalla letteratura sono organizzate in sei categorie e ponderate mediante l'Analytic Hierarchy Process sulla base dei giudizi di quattro esperti. Il DSS, implementato come pipeline Python riutilizzabile operante su una cartella di lavoro Excel strutturata, è applicato a un caso di eolico offshore incentrato su generatori a magneti permanenti e recupero di NdFeB, considerando una configurazione di riferimento e due scenari di mercato dei CRM. Il CBM orientato al prodotto ottiene il punteggio più alto nella configurazione di riferimento e nello scenario Low, mentre nello scenario High il CBM orientato all'uso ottiene il punteggio più alto. I risultati mostrano che la raccomandazione dipende sia dalle condizioni di mercato dei CRM sia dalle priorità del decisore.

Circular Business Models for Critical Raw Material Recovery: A Scenario-Based DSS with an Application in the Offshore Wind Sector

Abstract

The transition toward renewable energy has sharply increased demand for Critical Raw Materials (CRMs), exposing manufacturers to supply risk and price volatility. Circular Business Models (CBMs) provide alternative ways of organising ownership, responsibility and end-of-life control over an asset, but product-, use- and result-oriented archetypes entail different trade-offs. The literature provides limited structured support for comparing these alternatives from the perspective of a focal OEM under changing CRM market conditions.

This thesis develops a scenario-based Decision Support System (DSS) for comparing the three archetypes. Following a Design Science Research approach, decision variables derived from the literature are organised into six categories and weighted through the Analytic Hierarchy Process using the judgements of four experts. The DSS, implemented as a reusable Python pipeline operating on a structured Excel workbook, is demonstrated through an offshore wind case centred on permanent-magnet generators and NdFeB recovery, under a baseline and two CRM market scenarios. The product-oriented CBM receives the highest score under the baseline and Low scenarios, whereas the use-oriented CBM narrowly obtains the highest score under the High scenario. The results show that the recommendation is conditional on both CRM market conditions and the priorities assigned by the OEM, which the framework keeps explicit.

Table of Contents

Sommario.....	I
Abstract.....	II
1. Introduction.....	1
2. Literature Review	4
2.1 Critical Raw Materials	5
2.1.1 Strategic Relevance and Supply Risk.....	5
2.1.2 CRM Recovery and Circularity Challenges	8
2.1.3 Wind Energy and NdFeB Magnets as an Illustrative CRM-Intensive Case.....	9
2.2 Circular Business Models	12
2.2.1 CBM Definitions and Pattern Typology	12
2.2.2 Why a CBM Is Circular: Narrowing, Slowing, and Closing Resource Loops ..	14
2.2.3 The Control Dimension: Product-Service Systems.....	15
2.2.4 The Technical Recovery Dimension: The R-Strategies and Their Structural Link to the CBM	16
2.2.5 Actor Networks and Ecosystem Coordination in CBM Implementation	19
2.2.6 CRM Circularity and Value Chain Position: B2B vs B2C	20
2.3 Drivers, Barriers, and Market Conditions Affecting CBM Adoption	20
2.3.1 Determinants, Drivers and Barriers to CBM Adoption	20
2.3.2 Implementation Barriers and the Need for Stakeholder Coordination: Toward an OEM-Centred View	21
2.3.3 Market Uncertainty and Commodity-Price Volatility in Circular Strategies	22
2.4 Decision Support Systems and Multi-Criteria Decision-Making in the Circular Economy	24
2.4.1 Foundations: What a DSS Is.....	24
2.4.2 Circularity Indicators and Assessment Frameworks, and Their Limits for Comparing CBM Alternatives	24
2.4.3 DSS Applications in the Circular Economy	25
2.4.4 Multi-Criteria Decision-Making and the Analytic Hierarchy Process	26
2.5 Research Gap and Positioning of the Study	27
3. Methodology	30
3.1 Overall Research Design: A Design Science Research Approach	30
3.2 Method for Factor Identification	32
3.2.1 The OEM as Focal Actor and the Common Analytical Lens across the Three Archetypes	33

3.2.2 Value Chain Mapping of Each CBM Archetype across the Product Life Cycle	33
3.2.3 Cross-Model Comparison: Structural and Context-Derived Decision Factors	34
3.2.4 A Literature-Derived Coding Grid: Deriving Six Macro-Categories	35
3.2.5 Consolidation of Decision Factors into the Final Set of Decision Variables	37
3.2.6 The Decision Variable×CBM Matrix: Literature-Based Ratings and Benefit/Cost Classification	39
3.3 The MCDA Decision Engine	40
3.3.1 AHP Hierarchy and the Rationale for a Two-Level Structure	40
3.3.2 Pairwise Comparison and Weight Elicitation	41
3.3.3 Consistency Check	43
3.3.4 CBM Performance Matrix and Numerical Operationalisation	44
3.3.5 Weighted Scoring and CBM Ranking	45
3.4 Scenario Analysis Approach	46
3.4.1 Definition of CRM Market Scenarios	46
3.4.2 Scenario-Based Sensitivity Procedure	47
3.5 Interview Design and Data Collection	48
3.5.1 Expert Review and Structured AHP Elicitation	49
3.5.2 Expert Selection and Sampling	50
4. Framework Development: From CBM Value-Chain Mapping to the Decision Support System	54
4.1 Value Chain / Actor Mapping of the Three CBM Archetypes	54
4.1.1 Product-Oriented CBM	55
4.1.2 Use-Oriented CBM	56
4.1.3 Result-Oriented CBM	56
4.2 Cross-Model Comparison: From Value Chains to Candidate Decision Factors	57
4.2.1 Structural Divergence across the Three Value Chains	57
4.2.2 Context and Scenario Conditions Beyond the Comparison	59
4.3 The Final Set of Decision Variables	61
4.3.1 Market	61
4.3.2 Economic/Business Case	62
4.3.3 Technical/Product	64
4.3.4 Organisational/Capabilities	65
4.3.5 Value Chain/Ecosystem	66
4.3.6 Regulatory/Institutional	67
4.4 Rating the Archetypes and Testing Scenario Sensitivity	69

4.4.1 Assigning Low/Medium/High Ratings and Benefit/Cost Classification across the Three Archetypes.....	69
4.4.2 Normalising the Ratings into the Baseline CBM Performance Matrix	73
4.4.3 Testing the Baseline against Low and High Critical-Raw-Material Market Scenarios	76
4.4.4 Trade-Off Profile of the Three Archetypes across Macro-Categories	79
4.5 Decision Support System Framework.....	80
4.5.1 DSS Functional Architecture and Information Flow	81
4.5.2 AHP Weighting Hierarchy	85
4.5.3 CBM Performance Matrix and Weighted Aggregation	86
4.5.4 Scenario Integration and DSS Output	87
4.5.5 Software Implementation	87
5. Case Study: Offshore Wind Sector and NdFeB Magnet Recovery	89
5.1 Case Definition and Decision Context.....	89
5.2 Expert Elicitation Outcomes and Criteria Weights	90
5.2.1 Consistency of Expert Judgements.....	91
5.2.2 Weights, Expert Agreement and Interpretation	92
5.3 Baseline CBM Ranking	95
5.4 CBM Ranking under Alternative CRM Market Conditions	99
5.4.1 Low CRM Market Scenario	100
5.4.2 High CRM Market Scenario.....	102
5.5 Sensitivity to the Weights of the Individual Experts.....	105
5.6 Discussion	106
6. Conclusions.....	110
References	114
Appendix A	122
Appendix B	128
B.1 Overview	128
B.2 Source Code Listing	134
B.3 Program Output	140

1. Introduction

The transition toward renewable energy and other low-carbon technologies has sharply increased demand for Critical Raw Materials (CRMs), while their supply remains concentrated in a small number of countries and exposed to significant geopolitical and price risk (Cimprich et al., 2023; Gaustad et al., 2018). The European Union has responded with binding domestic targets for extraction, processing and recycling under the Critical Raw Materials Act, yet current end-of-life recycling rates remain far below these targets (Baldassarre & Carrara, 2025; van Gaalen & Slootweg, 2025). Overcoming this requires a shift toward circularity, but the barriers involved are not only technical but also managerial (Babbitt et al., 2021; Cimprich et al., 2023; van Gaalen & Slootweg, 2025). Offshore wind provides a particularly relevant setting for this problem. Direct-drive permanent magnet generators, which are widely used in offshore applications, rely on NdFeB permanent magnets and therefore create a direct link between wind-turbine deployment and demand for rare-earth magnet materials (Carrara et al., 2025). At the same time, the recovery of rare-earth magnets at end-of-life remains limited: fewer than five per cent of post-consumer rare-earth magnets are currently recycled worldwide (Chinwego et al., 2025; Rizos et al., 2024). The literature identifies a combination of organisational, market, supply-chain, technical and economic barriers that constrain the development of more effective recovery and recycling systems (Chinwego et al., 2025).

Circular Business Models (CBMs) provide a way of reconfiguring how ownership, responsibility and end-of-life control are allocated over an asset. The relative suitability of the CBM archetypes, product-, use- and result-oriented, is context-dependent, as they entail different trade-offs in terms of risk, capital requirements, operational responsibilities and recovery-related control (Reim et al., 2015; Tukker, 2004). Empirical reviews of CBM adoption identify market, organisational, technical, value-chain and regulatory conditions that jointly affect their implementation (Assmann et al., 2023). In a CRM-intensive context, these considerations are further affected by market conditions, since changes in CRM prices and supply conditions can alter the relative economic attractiveness of alternative configurations, while remaining largely outside the control of the OEM (Gaustad et al., 2018). This thesis therefore treats the OEM as the focal decision-maker, given its central position in the value chain and its influence over product design and end-of-life recovery (Cimprich et al., 2023).

Existing research addresses important parts of this decision problem, but the relevant perspectives remain only partially integrated. Criticality-assessment methods provide structured approaches to identifying material-related risks but offer more limited support for linking these assessments to the subsequent selection of mitigation strategies (Otterbach & Fröhling, 2025). Similarly, CBM-selection and multi-criteria decision-making approaches provide established bases for comparing alternatives but offer limited structured support for comparing CBM archetypes from the perspective of a focal OEM while explicitly accounting for changing CRM market conditions. The literature therefore provides substantial knowledge on the individual components of the problem, but less support for addressing them jointly within a single managerial decision.

This thesis addresses this problem by developing a scenario-based Decision Support System (DSS) for the comparative evaluation of the three CBM archetypes from an OEM perspective. The decision variables are derived from the literature on CBM configuration and implementation conditions and organised into a structured decision framework, whose relative importance is established through the Analytic Hierarchy Process (AHP) using the judgements of four experts. The resulting architecture is applied under alternative CRM market scenarios to examine how changes in market conditions affect the relative attractiveness of the CBM archetypes. The demonstration focuses specifically on permanent-magnet generators and NdFeB recovery in offshore wind. The technical feasibility of magnet recovery is treated as a given condition of the case rather than as an object of assessment, as the study focuses on the managerial choice of CBM configuration rather than on the optimisation of the recovery process itself. The research is guided by three questions:

1. *Which factors should inform an OEM's choice among Circular Business Models for the recovery of Critical Raw Materials?*
2. *How can these factors be translated into a managerial Decision Support System?*
3. *How does market uncertainty affect the resulting recommendation, and how can it be represented through explicit scenarios?*

Because the goal is to construct and evaluate a decision-support artefact rather than to explain or predict OEM behaviour, the research follows a Design Science Research (DSR) approach (Hevner et al., 2004).

The research presented in this thesis was initiated at the University of Genoa in March 2026 and completed at the University of Strathclyde, Glasgow, in September 2026, during a three-month Erasmus+ Traineeship at the Department of Design, Manufacturing and Engineering Management (DMEM), Faculty of Engineering. The work was conducted mainly in the Leonardo Suite, DMEM's postgraduate research community space, alongside PhD candidates and visiting researchers. The stay also provided a particularly relevant setting for the case study, given Scotland's prominent role in the offshore wind sector, and facilitated contact with the experts involved in the study.

Chapter 2 reviews the literature on critical raw materials, circular business models and existing decision-support approaches, establishing the problem and research gap that motivate RQ1. Chapter 3 sets out the methodology for deriving, organising and weighting the decision variables, including the design of the expert interviews, and for examining how the resulting recommendation changes under different CRM market scenarios. Chapter 4 consolidates these variables into the final decision-variable set and formalises the AHP architecture and its software implementation, so that RQ1 is addressed jointly across Chapters 2 and 4 and RQ2 primarily by Chapters 3 and 4. Chapter 5 applies the framework to an illustrative offshore wind case in which an OEM evaluates the three CBM archetypes in the context of permanent-magnet generators and NdFeB recovery. It reports the weights obtained from the experts and assesses the case under a baseline configuration and two alternative CRM market scenarios, allowing the study to examine how changes in CRM market conditions affect the relative attractiveness of the alternatives and the resulting recommendation. RQ3 is therefore addressed through the scenario-based application of the framework within this case. The chapter closes with a discussion of the results in light of the

literature and with general indications for the use of the DSS. Chapter 6 concludes by summarising the answers to the research questions, the contribution of the work, its limitations and directions for future research.

2. Literature Review

The global transition towards sustainable technologies has elevated the strategic importance of Critical Raw Materials (CRMs), while simultaneously exposing manufacturers to unprecedented supply chain vulnerabilities and price volatility. Addressing these risks requires a definitive shift from linear resource consumption to a circular economy. However, implementing circularity for CRMs is a multifaceted challenge: it is not just a technical or chemical problem, but a deeply organisational and managerial one. To fully capture the dimensions of this challenge, this chapter synthesises four distinct but interconnected bodies of literature, establishing the theoretical foundation for the decision-support framework developed in this thesis.

The chapter begins by exploring the literature on Critical Raw Materials (Section 2.1). It examines their strategic relevance, the resulting supply risks for downstream firms, and the specific design and systemic barriers that complicate their end-of-life recovery. To contextualise these challenges, NdFeB permanent magnets in offshore wind energy are introduced as an illustrative case, highlighting how CRM recovery often exposes organisational bottlenecks rather than just technical limitations.

Building on this context, Section 2.2 examines Circular Business Models (CBMs) as the structural mechanism for implementing circularity. This section reviews established CBM typologies and focuses on how different models, particularly Product-Service Systems, determine the level of control an Original Equipment Manufacturer (OEM) retains over its products. It explains how this dimension of control structurally dictates which end-of-life recovery strategies are viable and how they depend on a wider ecosystem of actors.

While the conceptual benefits of CBMs are well documented, their practical implementation involves significant challenges. Therefore, Section 2.3 evaluates the empirical record on the drivers, barriers, and market conditions that shape CBM adoption. This section narrows the focus from general determinants towards the specific perspective of the OEM, emphasising how market uncertainty and commodity-price volatility complicate the transition to circular strategies.

Navigating these complex, multi-stakeholder, and highly uncertain environments requires robust analytical tools. Section 2.4 reviews the analytical apparatus available to support these managerial choices, specifically Decision Support Systems (DSS) and Multi-Criteria Decision-Making (MCDM) methods. It evaluates how far existing circularity indicators and assessment frameworks have utilised these tools and identifies their limitations in comparing alternative business models.

Finally, Section 2.5 draws these four theoretical strands together. By demonstrating the lack of frameworks that combine CBM selection, market uncertainty, and a specific decision-maker's perspective, this concluding section formalises the research gap that this thesis seeks to address.

2.1 Critical Raw Materials

2.1.1 Strategic Relevance and Supply Risk

Critical Raw Materials occupy an increasingly strategic and contested position within the technological frameworks that the European Union relies on to meet its climate and industrial policy objectives, ranging from renewable energy generation to electric mobility and digital infrastructure (Baldassarre & Carrara, 2025). However, it is not an intrinsic physical property of a material, rather, it is a dynamic construct resulting from the systematic evaluation of two key dimensions: economic importance to end-use industrial sectors and the risk associated with its supply (Cimprich et al., 2023; Petrova, 2023). Instead of evaluating resource scarcity in isolation, assessments map candidate materials onto a two-dimensional criticality matrix to identify strategic vulnerabilities rather than ranking materials by physical scarcity alone (Gaustad et al., 2018; Petrova, 2023).

Fig. 2.1 sets out the recurring characteristics across these analytical assessments, demonstrating their compounding nature. High economic importance is precisely what elevates a geographically concentrated and geologically constrained supply from a minor operational bottleneck into a profound strategic vulnerability (Gaustad et al., 2018). This convergence of deep industrial reliance with a highly vulnerable resource base drives both corporate actors and policymakers to address the supply security of these materials with utmost urgency (Cimprich et al., 2023; Gaustad et al., 2018).




Fig. 2.1 - Basic conceptual characteristics of Critical Raw Materials (Petrova, 2023).

The classification of these materials is inherently dynamic and has undergone several structural revisions to reflect evolving industrial demands and geopolitical realities (Hool et al., 2024; Petrova, 2023). An example is the integration of copper and nickel, which have been designated as strategic raw materials under recent European policy initiatives (Hool et al., 2024). Although these metals do not strictly cross the traditional quantitative thresholds for supply risk due to their diversified global markets, they have been appended to the comprehensive critical raw materials framework in recognition of their indispensable and

irreplaceable role in decarbonisation technologies, thereby supporting targeted recycling and capacity-building initiatives (Hool et al., 2024; van Gaalen & Slootweg, 2025).

A similar analytical logic is adopted internationally, where jurisdictions customise criticality assessments to their specific socio-economic and military needs (Gaustad et al., 2018). For instance, the US Department of Energy evaluates materials based on their importance to clean energy technologies, while the Defense Logistics Agency focuses on materials essential to national security and defence, resulting in overlapping yet distinct resource priority rankings (Gaustad et al., 2018). This continuous evolution of international classifications underscores the core premise of this chapter: criticality is not a static inventory, but a moving target continuously shaped by technological trajectories, market dynamics, and shifting geopolitical landscapes (Gaustad et al., 2018; Hool et al., 2024).

The extraction and processing of most Critical Raw Materials are heavily concentrated in a small number of countries, creating structural vulnerabilities whereby supply shocks do not remain confined to the mining stage (Cimprich et al., 2023). When an upstream disruption occurs, the resulting price volatility and delivery bottlenecks propagate along the value chain, tending to affect firms positioned further downstream more severely than the initial point of origin would suggest an amplification effect comparable to the bullwhip phenomenon in supply chain management (Gaustad et al., 2018).

Corporate practices reflect a growing strategic awareness of this exposure. For instance, General Electric's ten-point criticality scoring matrix and Rolls-Royce's five-by-five risk matrix both integrate purchase volumes, historic price volatility, and substitution difficulty into a single score used to prioritise corporate mitigation actions (Gaustad et al., 2018). Concurrently, market concentration metrics such as the Herfindahl-Hirschman Index (HHI) are employed to track supply market oligopolies (Gaustad et al., 2018). At the institutional level, this metric is further combined with worldwide governance indicators to formally quantify supply risks under regulatory frameworks (von Rennenberg et al., 2024).

These risk assessment exercises are driven by high-stakes financial realities. Indeed, the 2006 price shock for rhenium is estimated to have cost General Electric an additional one billion dollars between 2006 and 2016 (Gaustad et al., 2018). This massive financial exposure prompted the firm to transition from a reactive posture to a proactive "revert, recover, recycle, and reduce" circular framework (Gaustad et al., 2018). By leveraging closed-loop business-to-business (B2B) maintenance and overhaul service contracts, the manufacturer successfully recovered rhenium from end-of-life turbine blades, significantly mitigating its exposure to primary supply and price fluctuations (Cimprich et al., 2023; Gaustad et al., 2018).

This critical exposure, however, is unevenly distributed across industrial sectors. Survey data collected from industry executives indicate that the automotive and infrastructure sectors report the highest proportion of products whose value or weight depends on scarce minerals and metals, with the renewable energy, utility, and power sectors also ranking among the most vulnerable downstream actors (Gaustad et al., 2018).




Fig. 2.2 - Sectors reporting the highest share of products by value and weight from scarce minerals and metals (Gaustad et al., 2018).

The European Union’s policy response has significantly evolved from a primarily diagnostic exercise, exemplified by the periodic updates of the Critical Raw Materials (CRM) list, towards more binding and assertive industrial measures (Baldassarre & Carrara, 2025; Hool et al., 2024). This regulatory shift is formally embodied in the Critical Raw Materials Act (CRMA), which frames domestic extraction, processing, and recycling as highly complementary levers designed to secure the EU’s open strategic autonomy (Baldassarre & Carrara, 2025).

To operationalise this strategy the CRMA, which entered into force in 2024, establishes ambitious quantitative benchmarks for Strategic Raw Materials (SRMs) to be achieved by 2030: domestic extraction must cover at least 10% of the EU’s annual consumption, domestic processing must reach at least 40%, and domestic recycling capacity must rise to a demanding 25% of annual demand, while reliance on any single third country must not exceed 65% for any strategic material (Carrara et al., 2025; van Gaalen & Sloomweg, 2025). This progressive escalation from the initial legislative proposals, which originally targeted a lower 15% recycling benchmark, highlights an active legislative effort to steer industrial practices towards self-reliance rather than just monitoring market exposures (Hool et al., 2024).

It is against this assertive policy backdrop that circularity re-enters the strategic discourse. In a conventional linear, “take-make-dispose” economic paradigm, manufacturing firms lack systemic mechanisms to recapture the material value of critical elements embedded in products once they are commercialised (van Gaalen & Sloomweg, 2025). Every unit of virgin critical material consumed must eventually be replaced by primary extraction, structurally reinforcing the very import dependencies and geopolitical vulnerabilities that threaten downstream industries (Bocken et al., 2016; Cimprich et al., 2023; van Gaalen & Sloomweg, 2025). Mitigating this exposure therefore requires a fundamental transition from linear resource flows towards circular systems that treat end-of-life products not as waste, but as strategic secondary reserves (Cimprich et al., 2023; van Gaalen & Sloomweg, 2025).

2.1.2 CRM Recovery and Circularity Challenges

Circularity is repeatedly proposed as the strategic response to the supply exposure and geopolitical risks associated with Critical Raw Materials (CRMs), yet the empirical recovery record for these elements remains critically modest. Current European Union end-of-life recycling input rates (EoL-RIR) for CRMs average only 7.3%, a figure that underscores how limited the industrial recapture of material value remains in practice and highlights the vast performance gap separating current activities from the Critical Raw Materials Act's (CRMA) 2030 target of 25% domestic recycling (van Gaalen & Slootweg, 2025).

The reasons underpinning this persistent gap are only partially chemical or metallurgical. While demand for critical metals, including cobalt, lithium, rare earth elements, and indium, continues to accelerate, their end-of-life recycling rates remain below 5% for most of these elements (Babbitt et al., 2021; Gaustad et al., 2018). This systemic failure is primarily rooted in product design: complex, high-technology products are rarely engineered to facilitate non-destructive disassembly, component upgrade, or efficient material separation at their end-of-life (Babbitt et al., 2021).

Addressing these vulnerabilities through circular strategies necessitates targeted design interventions that systematically narrow, slow, and close resource loops across the technology life cycle (Bocken et al., 2016; Babbitt et al., 2021). However, design-level modifications are insufficient in isolation, they require parallel systemic and institutional support mechanisms to establish the trust, transparency, and reverse logistics channels upon which second-life material markets depend (Babbitt et al., 2021; Gaustad et al., 2018). This is an important conceptual qualification: recycling technology itself is rarely the sole binding constraint. Because linear supply chains leave critical materials highly dispersed in minuscule quantities within complex alloys and multi-material components, such as indium in liquid crystal displays or rare earths in micro-electronics, even technically feasible recovery processes are frequently economically unviable without a fundamental restructuring of the surrounding business model (Babbitt et al., 2021; van Gaalen & Slootweg, 2025).

Shifting from transaction-based sales towards Product-Service Systems (PSS) and “product-as-a-service” arrangements represents a primary operational pathway to overcome this economic barrier (Bocken et al., 2016; Gaustad et al., 2018). By retaining ownership of the physical asset throughout its operational life, manufacturers gain a direct economic incentive to design for modularity, longevity, and ease of repair (Bocken et al., 2016). This ownership retention allows original equipment manufacturers (OEMs) to internalise the residual value of the high-grade critical materials contained within decommissioned assets, transforming what would traditionally be a costly waste-disposal liability into a highly valuable, closed-loop secondary resource stream (Cimprich et al., 2023; Gaustad et al., 2018).

Achieving the ambitious recycling benchmarks set by the CRMA is heavily anticipated to depend on complementary regulatory and technological instruments currently being deployed, most notably Extended Producer Responsibility (EPR) schemes and Digital Product Passports (DPPs) (Babbitt et al., 2021; van Gaalen & Slootweg, 2025). While individual EPR frameworks legally obligate manufacturers to take back end-of-life products, thereby forcing the internalisation of recovery costs (Gaustad et al., 2018), Digital Product Passports provide the necessary technical infrastructure by utilising material-tracking

technologies, blockchains, and IoT to store and share reliable, life-cycle-wide material composition data (Babbitt et al., 2021; van Gaalen & Slootweg, 2025). Together, these instruments aim to transform end-of-life recovery from a voluntary, market-conforming decision into a traceable, standardised, and enforceable obligation (van Gaalen & Slootweg, 2025).

Synthesising these perspectives reveals that the primary bottleneck restricting CRM circularity is as much managerial and structural as it is technological (Babbitt et al., 2021; Cimprich et al., 2023; van Gaalen & Slootweg, 2025). The specific circular business model an industrial actor adopts, alongside its positioning within the value chain (B2B vs. B2C), ultimately dictates whether a technically feasible recovery pathway can be translated into an economically viable industrial practice (Cimprich et al., 2023; Gaustad et al., 2018).

2.1.3 Wind Energy and NdFeB Magnets as an Illustrative CRM-Intensive Case

Not every technology intensive in Critical Raw Materials (CRMs) has received equal attention within the circular economy paradigm. While circularity research and innovation (R&I) and public funding have heavily concentrated on lithium-ion batteries, neodymium-iron-boron (NdFeB) permanent magnets remain comparatively under-researched, despite their strategic importance for wind energy, electric mobility, and digital infrastructure (Baldassarre & Carrara, 2025). This asymmetry represents a key research gap and a significant opportunity for industrial analysis (Baldassarre & Carrara, 2025).

The scale of the EU's material dependency is substantial. Technical assessments of the European wind sector show that offshore turbines, especially gearless, direct-drive Permanent Magnet Synchronous Generators (PMSGs), which are increasingly favoured for their mechanical reliability in harsh marine environments, rely on NdFeB magnets far more intensively per installed megawatt than their onshore, geared counterparts (Carrara et al., 2025). While onshore turbines traditionally utilise geared induction generators requiring little to no rare earths, modern offshore direct-drive models typically incorporate approximately 600 to 1,000 kg of NdFeB magnets per megawatt (Carrara et al., 2025). The rapid expansion of offshore wind capacity amplifies the EU's import reliance on a highly concentrated supply chain, where processing and magnet production are dominated by a single third country (Carrara et al., 2025).

This strategic vulnerability highlights that the primary barrier to closing rare-earth loops is a complex challenge of industrial pre-treatment and value-chain coordination (Chinwego et al., 2025; Couto et al., 2024). Laboratory-scale work on the pre-treatment of spent NdFeB magnets shows that physical recovery is well-characterised but remains operationally slow and highly labour-intensive (Couto et al., 2024). For instance, manual disassembly of a single hard-disk drive (HDD) to retrieve its voice coil motor magnets can take ten to fifteen minutes (Couto et al., 2024).

To overcome this manual bottleneck, advanced automated dismantling systems have been piloted, such as the rotating drum disassembly technology developed by Hitachi Group, which scales recovery performance from 10-12 HDDs per hour to approximately 140 HDDs per worker-hour (Cimprich et al., 2023). However, wind turbine generators introduce application-specific design constraints that resist standardised automation: magnets are tightly embedded, heavily glued, and coated with protective layers that vary widely by manufacturer (Couto et al., 2024).

Because offshore wind turbines operate under severe marine environments, their permanent magnets require highly durable anti-corrosion coatings (Couto et al., 2024). The presence of these protective layers represents a major barrier for direct recycling processes, as any residual coating material can contaminate the recovered magnetic fractions and severely degrade the performance of the remanufactured magnets (Couto et al., 2024). Bypassing this technical hurdle requires additional chemical or thermal pre-treatment stages to safely strip the coatings, which inevitably adds operational complexity and cost to the recovery chain (Couto et al., 2024).

Beyond these technical and chemical intricacies, the commercial viability of rare-earth magnet recycling is heavily constrained by structural market barriers (Chinwego et al., 2025; Rizos et al., 2024). Currently, fewer than 5% of post-consumer rare-earth magnets are recycled globally (Chinwego et al., 2025). This extremely low rate is not due to process infeasibility, but to the highly fragmented nature of end-of-life collection and the delayed retirement of key assets. Wind turbines and electric vehicles have not yet reached decommissioning in the large, predictable volumes required to sustain a stable commercial recycling infrastructure (Chinwego et al., 2025; Rizos et al., 2024).

Furthermore, the profitability of recycling operations is undermined by the high capital costs of recovery processes, competition with cheap primary magnets, and the extreme price volatility of virgin rare-earth elements on global markets (Rizos et al., 2024). Because the primary revenue stream for recycling firms relies on the direct business-to-business (B2B) sale of recovered neodymium-praseodymium (NdPr) alloys to magnet manufacturers, any commercial operation remains highly vulnerable to feedstock instability (Chinwego et al., 2025).

Establishing a resilient circular economy for NdFeB magnets depends on robust stakeholder collaboration, early customer discovery, and a workable corporate cost-and-revenue structure that mitigates market uncertainties (Chinwego et al., 2025). The primary constraints restricting NdFeB recovery are fundamentally managerial and systemic (Chinwego et al., 2025; Rizos et al., 2024). The specific circular business model an industrial actor adopts, and the mechanism it uses to retain control over end-of-life assets, ultimately determines whether a technically proven chemical extraction pathway can survive as an economically viable industrial practice.

The wind energy sector illustrates a closely related pattern of circular complexity at the industrial level. End-of-life wind turbines, and their composite blades in particular, are projected to generate a sharply accelerating volume of waste as the rapid global expansion of wind capacity drives up decommissioning rates (Tyurkay et al., 2024). According to institutional projections by the European Commission, cumulative waste from decommissioned turbines within the European Union is expected to exceed 4.7 million tonnes by 2030 (Carrara et al., 2025). Within this stream, composite wind turbine blades represent a major bottleneck, generating approximately 60,000 tonnes of complex waste in 2023, with volumes forecast to skyrocket to 400,000 tonnes annually by 2040 (Carrara et al., 2025).

This waste management challenge is heavily compounded by a fragmented European regulatory landscape: while certain member states have instituted outright bans on the landfilling of composite products (such as Germany and the Netherlands), others continue to rely on landfill taxes and gate fees of varying severity (Beauson et al., 2022; Tyurkay et

al., 2024). Every strategic disposition choice, ranging from continued operation versus decommissioning, to structural reuse versus mechanical processing, and advanced recycling versus final disposal, must be evaluated through a multi-dimensional lens that simultaneously accounts for economic viability, technical feasibility, legislative constraints, and environmental impacts (Beauson et al., 2022).

The consequences of failing to plan for these end-of-life waste streams extend well beyond environmental degradation. A renewable energy sector perceived as accumulating massive volumes of unmanaged composite waste risks severely damaging public acceptance and political trust, thereby undermining the core narrative of wind energy as a genuinely clean and sustainable technology (Beauson et al., 2022). This paradox, often referred to as the “green on green” conflict, highlights how the environmental credibility of the energy transition is compromised when downstream waste management is neglected (Beauson et al., 2022). Indeed, without proactive and standardised circular strategies, the rapid deployment of low-carbon technologies may fail to deliver its full socio-environmental potential, as the environmental benefits of green electricity are offset by the accumulation of unmanaged composite waste (Baldassarre & Carrara, 2025).

Recent literature characterises this systemic bottleneck as a circular dilemma, where the lack of large-scale recycling infrastructure limits the reliable supply of high-quality secondary materials, while the absence of robust market demand for recycled composites simultaneously discourages the private capital investments required to build such recycling facilities (Ibarra et al., 2026). This structural challenge cannot be addressed by technical innovations alone (Ibarra et al., 2026; Mendoza et al., 2022). Shifting the industrial focus from a narrow technological debate to a systemic and managerial perspective is essential to design and coordinate a viable, self-sustaining circular business ecosystem (Ibarra et al., 2026; Mendoza et al., 2022). Overcoming these systemic deadlocks requires robust evaluation frameworks capable of assessing alternative circular and sustainable business configurations, enabling manufacturers to systematically align emerging recovery technologies with progressive regulatory pressures and volatile secondary material markets (Ibarra et al., 2026).

When integrated, the empirical evidence from both rare-earth magnet recovery and wind turbine decommissioning converges on a singular conclusion. Determining whether critical material circularity is achieved depends not only on technical feasibility, but on identifying which industrial actors are strategically positioned, operationally organised, and economically incentivised to execute the recovery under viable business arrangements. This is fundamentally a question of business model innovation and inter-organisational ecosystem coordination. It is this coordination challenge that the remainder of this chapter addresses. Of the two illustrative streams introduced here, the permanent-magnet case is the one taken up in Chapter 5, where it provides the demonstration context for the framework developed in Chapter 4; the wind-turbine blade stream is discussed as a parallel example of the same managerial pattern and remains outside the scope of the case.

2.2 Circular Business Models

2.2.1 CBM Definitions and Pattern Typology

At its core, a Circular Business Model (CBM) represents a fundamental strategic shift from the traditional linear “take-make-dispose” industrial paradigm (Bocken et al., 2016; Geissdoerfer et al., 2023). Rather than focusing on throughput maximisation, a CBM is defined as the organisational logic by which an enterprise creates, delivers, and captures value by systematically cycling, extending, intensifying, and dematerialising material and energy loops, thereby reducing resource inputs and minimising waste and emissions leakage (Geissdoerfer et al., 2023). This conceptualisation bridges the gap between macroeconomic circular principles and firm-level strategic execution, positioning business model innovation as the primary driver of circular transitions (Geissdoerfer et al., 2023; Lewandowski, 2016).

The necessity of explicitly defining CBMs as a strategic vehicle is underscored by systematic literature reviews (Kirchherr et al., 2017). An analysis of 114 definitions of the circular economy revealed that only 11% explicitly identified business models as an enabler of circularity (Kirchherr et al., 2017). Instead, the broader literature historically conceptualised circularity through the lens of waste hierarchies and physical material flows, with a strong focus on the classical 3R or 4R principles (reduce, reuse, recycle, recover) (Kirchherr et al., 2017). While recycling (79%) and reuse (74-75%) were extensively mentioned in these definitions, the higher-level systemic shift and the active role of business models as drivers of these cycles were largely neglected, a gap that has necessitated the rapid development of a dedicated, managerially focused CBM literature (Kirchherr et al., 2017; Lewandowski, 2016).

Over the past decade, this applied literature has increasingly converged on a consolidated set of recurring business model design options and patterns (Lüdeke-Freund et al., 2019). Through a morphological analysis of 26 distinct CBM concepts identified in the literature, Lüdeke-Freund et al. (2019) consolidated these diverse approaches into six major patterns with the potential to close or slow resource loops: (1) repair and maintenance, (2) reuse and redistribution, (3) refurbishment and remanufacturing, (4) recycling, (5) cascading and repurposing, and (6) organic feedstock business models. To structure this taxonomy, each pattern is systematically evaluated across four fundamental business model dimensions: the value proposition (products and services), value delivery (target customers and delivery processes), value creation (partners and processes), and value capture (revenues and costs) (Lüdeke-Freund et al., 2019).

A central insight of this taxonomy is that circular business models require a fundamentally different financial and operational logic compared to linear systems (Lüdeke-Freund et al., 2019). While linear supply chains are optimised around “value added” along a unidirectional, forward resource flow, circular models are built on “value retained” across multiple lifecycles (Lüdeke-Freund et al., 2019). Because traditional economic assessments are optimised for high-volume forward sales, they frequently undervalue and fail to capture the long-term, multi-cycle economic potential of circular strategies, which is why alternative, systemic design frameworks are required (Lewandowski, 2016; Lüdeke-Freund et al., 2019).

To practically operationalise these four dimensions for business design, the literature has adapted traditional strategic tools like the Business Model Canvas (Lewandowski, 2016; Daou et al., 2020). Two distinct frameworks have emerged to guide manufacturers in this transition. Specifically, Lewandowski (2016) developed the Circular Business Model Canvas (CBMC), extending the traditional nine-block structure with two circular-specific components: the take-back system, which structures reverse logistics, collection, and customer relations for returned assets, and adoption factors, which capture the organisational capabilities and external PEST (political, economic, social, technological) conditions required to enable the transition. Alternatively, Daou et al. (2020) proposed the Ecocanvas, a circular design framework specifically engineered to integrate life-cycle thinking, social foresight, and environmental impact assessments directly into the business planning stage. Both of these strategic tools provide structured methodologies to ensure that the technical potential of resource recovery is backed by a viable, profit-oriented business logic (Daou et al., 2020; Lewandowski, 2016).

In operationalising these frameworks, research highlights specific implementation challenges that go beyond simple layout design. For instance, Lewandowski (2016) formalises what he terms the “triple fit challenge”, arguing that a circular business model achieves commercial viability only through the simultaneous alignment of three distinct, interdependent fits: the alignment between the circular value proposition (incorporating the take-back system) and the target customer segments; the balance within the financial engine, ensuring that the cost structure and revenue streams are sustainable and profitable; and the congruence between the organisational changes implemented towards circularity and the internal and external adoption factors that might otherwise constrain the process.

This requirement for systematic alignment is further illustrated by the empirical application of the Ecocanvas (Daou et al., 2020). When testing this foresight-driven tool within small and medium-sized wineries, researchers found that integrating environmental, social, and economic-legal dimensions directly into business planning was highly effective in surfacing external vulnerabilities, such as resource scarcity, price volatility, and tightening waste regulations (Daou et al., 2020). These surfaced risks are identical in nature to the geopolitical and supply exposures that this thesis addresses for Critical Raw Materials. These canvas-based extensions confirm that a circular business model cannot be evaluated on its value proposition alone; rather, its financial viability, reverse logistics feasibility, and organisational adoption capacity are constitutive of whether a circular strategy is genuinely viable for a firm, rather than an afterthought (Lüdeke-Freund et al., 2019).

For the purposes of this thesis, however, not all of these business model dimensions carry equal analytical weight. Because the core research questions concern the OEM’s choice among circular business models for the recovery of Critical Raw Materials from CRM-intensive industrial assets, the analysis that follows concentrates on the two dimensions most directly implicated in enabling this recovery: the ownership and control dimension (who retains custody of the asset over its life cycle) and the technical recovery dimension (how the physical materials are recaptured at end-of-life) (Assmann et al., 2023; Pollard et al., 2023).

While the broader value proposition and partner network dimensions documented in the literature remain highly relevant to a complete account of circularity and resurface where the

decision-support framework of Chapter 4 evaluates cost structures and revenue-stream trade-offs, they are not treated as independent axes of analysis here. This deliberate boundary selection is theoretically supported by empirical reviews demonstrating that securing end-of-life product control and establishing viable reverse logistics represent the two most critical organisational and operational bottlenecks for manufacturers transitioning towards high-value material recovery (Assmann et al., 2023; Pollard et al., 2023). Accordingly, the next sections develop the theoretical foundations for these two focus dimensions.

2.2.2 Why a CBM Is Circular: Narrowing, Slowing, and Closing Resource Loops

Before separating organisational control from technical recovery, it is essential to establish the theoretical foundations of circularity within business models. Bocken et al. (2016) address this fundamental question by distinguishing three primary mechanisms through which resource loops are managed, building upon an early conceptualisation of a performance economy where value is generated by maintaining and extending the utility of existing goods rather than maximising production volume. These mechanisms are defined as: narrowing loops, which focus on resource efficiency by reducing material and energy inputs per product without necessarily altering product lifespan; slowing loops, which extend product utilisation periods through durable design, maintenance, repair, and remanufacturing; and closing loops, which capture post-use materials to recycle them back into new production cycles once a product can no longer be kept in service. This terminology does more than classify corporate practices, it distinguishes strategies operating at different levels of organisational activity. Specifically, Bocken et al. (2016) treat circular product design strategies and circular business model strategies as parallel but distinct layers. The former governs the physical configuration of the product for durability, ease of disassembly, or standardisation, whereas the latter dictates how the firm delivers value and captures revenue.

When applied to critical raw materials (CRMs), this dichotomy becomes essential for mitigating supply-chain vulnerabilities and resource risks. As argued by Babbitt et al. (2021), mitigating CRM supply risks across the technology life cycle requires targeted circular design interventions, such as modularity or component standardisation, yet physical design changes remain insufficient without aligned business model configurations. Within Bocken et al.'s (2016) classification, the primary business model strategy supporting loop-slowness pathways is the access-and-performance model, where firms deliver capability or utility rather than transferring product ownership. This archetype belongs to the broader family of Product-Service Systems (PSS). Retaining ownership under a PSS framework provides original equipment manufacturers (OEMs) with both the operational information and the financial incentives to prolong product lifespans through maintenance and remanufacturing before eventually closing material loops through recycling. Fig. 2.3 sets out the full framework, distinguishing circular product design strategies from circular business model strategies across both the slowing and the closing categories (Bocken et al., 2016). Access-and-performance arrangements, and Product-Service Systems more broadly, are therefore not a technical recovery option on the same footing as remanufacturing or recycling; they belong to the business model layer that determines whether those technical options become available to the firm at all. Section 2.2.3 develops this control dimension in full, drawing on Tukker's (2004) wider taxonomy of product-service systems, of which the access-and-performance model is one specific instance.



Fig. 2.3 - Circular product design and business model strategies, organised by their effect on resource loops (slowing vs. closing) (Bocken et al., 2016).

2.2.3 The Control Dimension: Product-Service Systems

The control dimension of a circular business model determines which actor retains control over the physical product and its constituent materials throughout the lifecycle. This dimension is best captured by Tukker's (2004) taxonomy of Product-Service Systems (PSS) which remains the foundational reference point in this literature. Tukker (2004) conceptualises a spectrum running from pure product sales to pure service delivery, classifying eight distinct PSS types into three macro-categories:

1. Product-oriented systems remain closest to conventional transactions, ownership of the physical asset is transferred to the customer, while the manufacturer provides supplementary services, such as maintenance, repair, or consultancy.
2. Use-oriented systems shift the focus from product ownership to product access, the provider retains asset ownership and sells its utility through leasing, renting, or pooling arrangements.
3. Result-oriented systems represent the furthest point on the spectrum, where the customer buys a specific outcome, unit of service, or functional result (e.g., pay-per-service-unit) without any physical ownership ever transferring.

Each macro-category implies a distinct allocation of ownership, financial risk, and operational control, directly conditioning how end-of-life recovery is managed. Fig. 2.4 illustrates this taxonomy, showing how Tukker's (2004) eight subcategories nest within the three PSS families, and clarifying how Bocken et al.'s (2016) access-and-performance model maps across use-oriented (access) and result-oriented (performance) boundaries. For the purposes of this thesis, these three PSS macro-categories constitute the alternative Circular Business Models compared for critical raw material (CRM) recovery from Chapter 4 onward.

This taxonomy did not develop in isolation. In parallel, a substantial body of literature on the "servitisation" of manufacturing was emerging, describing a similar transition from standalone physical products towards integrated product-service offerings. However, this transition was primarily examined through a manufacturing-strategy lens rather than from a sustainability perspective, which contributed to limited integration between the two research streams despite their substantial thematic overlap (Baines et al., 2009). The convergence of

these distinct academic streams around comparable forms of value creation and delivery highlights the importance of lifecycle control and asset ownership as strategic levers for industrial transition. This theoretical intersection also supports the broader terminology adopted in this thesis, which uses Circular Business Models as an overarching concept rather than referring exclusively to Product-Service Systems.




Fig. 2.4 - Main and subcategories of Product-Service Systems (Tukker, 2004).

This taxonomy should not, however, be read as an implicit hierarchy in which more service-centric models automatically possess superior circularity or sustainability. In his seminal analysis, Tukker (2004) demonstrated that different PSS archetypes exhibit highly divergent environmental and economic profiles. Most PSS configurations yield only marginal ecological improvements, with the most radical resource-efficiency potential restricted to result-oriented models, which conversely face the most formidable implementation barriers and risk-sharing challenges. A decade later, a systematic review of 278 papers by Tukker (2015) confirmed that while the literature had matured considerably in mapping the strategic benefits of PSS, their actual sustainability contributions remain highly context-dependent, and systemic implementation challenges persist. This is particularly evident in the divergent dynamics between business-to-business (B2B) and business-to-consumer (B2C) markets, as well as the risk of suboptimal user care in non-ownership models.

This dual finding directly informs the operationalisation of the control dimension in this study. Rather than prescribing extreme servitisation as a default optimal path, the choice among product-, use-, and result-oriented configurations is treated as an empirical inquiry. Because the viability of these business models depends heavily on specific regulatory, technical, and market parameters, this thesis addresses CBM selection case-by-case. By utilising the multi-criteria decision-support framework constructed in Chapters 3 to 5, this study avoids taxonomic assumptions, evaluating the performance of each alternative against the concrete operational realities of critical raw material recovery in the offshore wind sector.

2.2.4 The Technical Recovery Dimension: The R-Strategies and Their Structural Link to the CBM

While the control dimension determines the organisational architecture of circular business models, the technical recovery dimension governs the physical transformation of products and materials once recaptured. The literature on closed-loop supply chains (CLSCs) and

reverse logistics provides the foundational taxonomy for these physical operations. Fleischmann et al. (1997) systematise product recovery options along a spectrum of increasing rework intensity, distinguishing between: direct reuse, which requires no physical modification; repair, which restores basic functionality; remanufacturing, which returns the asset to the OEM's original performance specifications while preserving product identity; and recycling, which recovers material value without conserving product structures.

Subsequent systematic reviews have refined this empirical landscape. Govindan et al. (2015), in a comprehensive review of 382 papers published between 2007 and 2013, categorised CLSC research across several structural dimensions, including decision levels, product types, and operational uncertainties. They observed that the choice of return channel has historically been treated in the literature as an operational given rather than a strategic decision driven by explicit economic-financial determinants.

To address this decision-making gap, narrower, sustainability-focused reviews have sought to structure the criteria behind product disposition. Govindan and Soleimani (2017) examined this transition, focusing on the specific decision-making gates for end-of-life products. They analysed a structured disposition framework which maps recovery options (repair, reconditioning, remanufacturing, recycling, and disposal) as distinct decision gates rather than a single aggregated recovery choice (Fig. 2.5).




Fig. 2.5 - Framework of factors influencing reverse logistics product disposition decisions (Govindan & Soleimani, 2017).

Empirical, firm-level evidence explains why the return channel is often treated as a structural constraint rather than a flexible strategic choice. In their survey of 89 ISO 14001-certified manufacturers, Khor and Udin (2013) demonstrated that a product's physical characteristics, specifically its design for disassembly, serve as the primary determinant for enabling high-value technical recovery pathways such as reconditioning, remanufacturing, and recycling. Conversely, general resource commitment and design for the environment were found to primarily influence lower-intensity options like repair and disposal, rather than acting as universal enablers of high-value recovery.

This discrepancy, between a well-established taxonomy of technical recovery options and the comparatively scarce explicit treatment of the strategic and economic-financial determinants of the choices between them, motivates the approach of this thesis. The framework developed from Chapter 4 onward is therefore designed as a strategic, multi-criteria decision-support model rather than a techno-economic one. Rather than optimising physical chemical yields, sizing processing equipment or modelling cash flows, this study treats the technical recovery processes, at the granularity established by Fleischmann et al. (1997) and verified by Govindan et al. (2015) and Govindan and Soleimani (2017), as given conditions against which the attractiveness of circular alternatives is compared from the focal OEM's perspective, with economic and financial considerations entering as weighted decision variables (Chapter 4).

What Section 2.2.2 and Section 2.2.3 suggest, and what this discussion makes explicit, is that organisational control and technical recovery are highly interdependent. Bocken et al. (2016) formalise this linkage specifically for Product-Service Systems: when an OEM retains ownership under an access-and-performance model, it inherently internalises both the operational costs and the end-of-life liabilities of the physical asset. This arrangement fundamentally realigns corporate incentives, discouraging planned obsolescence and instead rewarding product configurations designed for durability, modularity, and ease of disassembly. Under this business model, continuous revenue streams from service delivery offset the upfront capital expenditures of maintenance and remanufacturing. By contrast, conventional linear sales lack this incentive structure, as the manufacturer relinquishes residual claims and lifecycle visibility at the point of sale. Recycling, at the closing end of the R-strategy hierarchy, frequently becomes the only viable recovery pathway once the product exits the OEM's line of sight. This mechanism aligns with the core idea that selling utility over time provides a continuous financial rationale to circulate products and components through loop-slowness pathways (repair and remanufacturing) before material recycling becomes unavoidable (Bocken et al., 2016). Accordingly, a central premise of this thesis is that the chosen CBM (product-, use-, or result-oriented) structurally conditions which point on the technical recovery hierarchy (repair, reconditioning, remanufacturing, or recycling) is economically feasible, rather than the two dimensions being independently combinable choices.

To empirically evaluate this structural link, this study focuses on neodymium-iron-boron (NdFeB) permanent magnets, whose physical recovery constraints are already well-characterised at the laboratory and pilot scale. As documented by Couto et al. (2024), retrieving these critical materials from end-of-life products involves a complex and highly variable process chain: manual disassembly, thermal demagnetisation, elemental characterisation, and selective chemical de-coating. Each of these steps is inherently constrained by labour intensity and the structural unpredictability of feedstocks. Because the technical and engineering feasibility of magnet recovery is not in question, the analytical contribution of this thesis is directed elsewhere. Specifically, this study investigates under which market and organisational conditions each CBM archetype through which the OEM can enable a given technical recovery option is comparatively preferable, a question Section 2.2.5 situates within stakeholder networks, and Section 2.3.3 addresses under conditions of commodity-price volatility and market uncertainty.

2.2.5 Actor Networks and Ecosystem Coordination in CBM Implementation

Analysing CBMs from the standpoint of a single focal firm, as done in Sections 2.2.3 and 2.2.4, provides a useful analytical simplification; however, it risks underestimating the degree to which circular business model (CBM) implementation depends on actors beyond the focal firm. Under a systemic lens, circularity is conceptualised as a property of a higher-order system rather than an attribute of an individual product or organisation (Konietzko et al., 2020a). Closing material loops typically requires the strategic alignment of business models among multiple stakeholders, including suppliers, customers, and regulators, presenting a complex coordination challenge that no single firm can resolve unilaterally (Konietzko et al., 2020a). This systemic view aligns with broader sustainability management theory, which argues that corporate environmental and economic performance depends on transforming conflicting stakeholder interests into shared, mutually reinforcing goals, rather than treating commercial and ecological objectives as inherently antagonistic (Hörisch et al., 2014).

Even when the intent to coordinate exists, the experimentation and implementation of circular configurations may not follow a fully linear or systematic process. Konietzko et al. (2020b) highlight the role of experimentation and effectual reasoning in circular business model innovation, whereby decisions can be shaped by the actors' available means, including who they are, what they know, and whom they know. This perspective suggests that some assumptions underlying circular configurations may remain implicit rather than being explicitly articulated and assessed. Against this background, the actor-mapping methodology introduced in Chapter 3 provides a structured way to make relevant actors, roles, and interdependencies explicit and therefore more amenable to systematic assessment.

Applied studies in the electrical and electronic equipment (EEE) sector operationalise this collaborative argument. Pollard et al. (2021) propose a five-layered Circular Economy Business Model Innovation (CEBMI) process framework. Developed through seven iterative workshops with EEE stakeholders, this framework systematically connects business strategy, a sector-specific CEBM Canvas, implementation challenges, regulatory policies, and micro-level indicators. This multi-actor approach demonstrates that transitioning from a conceptual commitment to circularity to an operational one requires structured, collaborative processes rather than isolated corporate decisions.

Similarly, a subsequent EEE sector case study applies a nine-block CEBM Canvas across eight concurrent focus groups (Pollard et al., 2023). This study reveals that Product-Service Systems (PSS) consistently emerge as the most attractive path to secure recurring revenue, primarily because they allow manufacturers to retain material ownership. Reverse logistics and predictive maintenance are identified as the core operational activities that make such material retention practically viable rather than merely aspirational (Pollard et al., 2023).

When read alongside the ecosystem-innovation principles, collaboration, experimentation, and platformisation, proposed by Konietzko et al. (2020a), these findings converge on a shared conclusion. Circular business model implementation is not an isolated organisational decision, but a complex coordination challenge that must be iteratively resolved with external actors.

Platformisation, understood as governing data flows and information-sharing arrangements, is particularly critical for critical raw material (CRM) recovery (Konietzko et al., 2020a). Indeed, tracking material location and product composition across supply chains represents a distinct organisational hurdle, independent of the physical recovery process itself. This

actor-centric, multi-stakeholder logic is precisely what the value chain and actor-mapping method of Section 3.2 formalises for the three CBM alternatives introduced in Section 2.2.3, providing the theoretical foundation for the general framework developed in Chapter 4.

2.2.6 CRM Circularity and Value Chain Position: B2B vs B2C

The coordination challenge described in the preceding section varies significantly across value-chain configurations. Business-to-business (B2B) settings are generally more conducive to circularity for critical raw materials (CRMs) than business-to-consumer (B2C) contexts (Cimprich et al., 2023). This advantage stems primarily from clearer property rights and lower transaction costs, which grant the focal firm tighter control over products at their end of life (Cimprich et al., 2023).

To evaluate these dynamics systematically, Cimprich et al. (2023) mapped circular strategies across three distinct states of matter: particles, parts, and products. Based on five industrial cases, their analysis reveals that closed-loop strategies, which contain material flows within the focal actor's system boundary, are what secure supply and mitigate price volatility for that actor. Conversely, open-loop recycling, where recovered materials cross the system boundary and exit the manufacturer's value chain, does comparatively little to reduce the focal firm's supply risk, even if technically successful. B2B arrangements, such as leasing, service contracts, and maintenance agreements, facilitate circularity by internalising operational costs that would otherwise fall on fragmented downstream actors. In contrast, B2C settings typically require robust external policy incentives (e.g., deposit-refund systems or trade-in credits) or significant technological breakthroughs to approximate comparable coordination efficiency (Cimprich et al., 2023).

Nevertheless, recovery in B2C markets is not entirely impossible. Empirical evidence from the Hitachi case study demonstrates how the recovery of NdFeB magnets from end-of-life computer hard-disk drives (a B2C product) was made economically viable through a combination of automated disassembly technology and government recycling subsidies (Cimprich et al., 2023). These mechanisms effectively substitute for the default ownership control typical of B2B relationships. Read against the control-dimension argument developed in Sections 2.2.2 to 2.2.4, this represents a natural test of the underlying claim: where ownership control is present (as in B2B), recovery is organisationally straightforward, where it is absent (as in B2C), recovery requires compensatory mechanisms, such as automated sorting or policy subsidies, to overcome coordination barriers.

Consequently, an original equipment manufacturer (OEM) occupying a central position in a B2B value chain, such as a wind turbine manufacturer coordinating component suppliers, installers, and dismantlers, is uniquely positioned to internalise the coordination challenges identified in Section 2.2.5. Rather than leaving material loops unmanaged across a fragmented network of independent actors, the OEM can leverage its structural position to lead circular implementation. Section 2.3.2 will explicitly justify why this study focuses on the OEM, rather than recyclers or regulators, as the central decision-making actor.

2.3 Drivers, Barriers, and Market Conditions Affecting CBM Adoption

2.3.1 Determinants, Drivers and Barriers to CBM Adoption

Two recent systematic reviews provide a comprehensive taxonomy of the factors enabling and constraining circular business model (CBM) adoption. Assmann et al. (2023) synthesise

the literature into 54 distinct determinants categorised across eight macro-dimensions, spanning cultural, regulatory, market, strategic, business case, value chain, technological, and knowledge conditions.

Complementing this, Geissdoerfer et al. (2023) identify 25 barriers and 10 drivers clustered into seven categories. Their work includes the perspective of incumbent firms, established manufacturers evaluating a circular transition within existing linear setups, rather than focusing solely on circular start-ups entering the market. This incumbent-firm framing is highly relevant to original equipment manufacturers (OEMs), whose barriers to CBM adoption are shaped as much by the friction of altering legacy linear operations as by any general deficit in technology or capital.

These structured reviews align with earlier empirical evidence. A large-scale survey conducted across the European Union reveals that, among the barriers ranked as most critical by businesses and policymakers, none were technological in nature. Instead, the most prominent hurdles were cultural, specifically, hesitant corporate mindsets and low consumer awareness and market-driven, primarily the uncompetitively low price of virgin materials relative to secondary alternatives (Kirchherr et al., 2018).

A complementary strand of literature traces these market and cultural barriers back to upstream regulatory shortcomings. Applying a fuzzy-DEMATEL causal analysis to nineteen policy-related barriers within an industrial case, Kazancoglu et al. (2021) demonstrate that market-based obstacles are frequently induced by insufficient regulatory intervention. This includes a lack of circular public procurement, inadequate financial incentives, and poor coordination between environmental and industrial policies. What superficially appears to be market failure is often a symptom of policy failure further upstream.

Consequently, these four sources converge on a core premise of this thesis: CBM adoption is constrained less by technological feasibility than by the economic viability of the business case and the stability of surrounding regulatory and market conditions. This fundamental uncertainty justifies the application of a scenario-based decision-support tool, which is uniquely suited to help firms commit capital under volatile external conditions.

2.3.2 Implementation Barriers and the Need for Stakeholder Coordination: Toward an OEM-Centred View

Among the determinants catalogued in Section 2.3.1, financial conditions are consistently identified as the most binding constraints. In a systematic review of 126 empirical studies on circular business model adoption, Hina et al. (2022) classify the obstacles faced by firms into internal and external categories. External barriers encompass supply-chain fragmentation, regulatory gaps, and consumer scepticism regarding recycled product quality. Internally, firms are primarily constrained by high upfront investment requirements, a shortage of specialised managerial expertise, and financial risks. An uncertain return on investment (ROI) and prohibitive initial capital costs are repeatedly singled out as the most limiting internal hurdles (Hina et al., 2022). These findings, which highlight economic viability rather than technological deficit as the primary barrier, motivate the focus of this research.

Because technology itself is rarely the limiting factor, the orchestration of the actors who collectively enable material recovery becomes the decisive variable. As established in Section 2.2.5, closing material loops requires the structural alignment and compatibility of

business models across multiple stakeholders. This systemic challenge cannot be resolved by a single firm acting unilaterally (Konietzko et al., 2020a). More broadly, sustainable value creation depends on aligning conflicting stakeholder interests into mutually reinforcing objectives rather than treating economic and ecological goals as antagonistic (Hörisch et al., 2014).

In a CRM recovery context, this alignment challenge is highly concrete. For example, a wind farm operator's objective to minimise decommissioning costs does not automatically align with a recycler's requirement for a stable, high-volume feedstock. These distinct goals only become compatible through a viable collaborative business model. Recovering NdFeB magnets from wind turbines depends precisely on this inter-firm coordination, as it requires reliable feedstock logistics, transparent information regarding magnet location and physical composition, and formal agreements between manufacturers, dismantlers, and recyclers that secure economic benefits for all participants (Chinwego et al., 2025).

However, as Section 2.2.6 demonstrated, the difficulty of resolving this coordination problem varies by value-chain configuration. Business-to-business (B2B) settings significantly reduce transactional and coordination costs compared to business-to-consumer (B2C) markets because asset ownership is concentrated and property rights are clearly defined among fewer, larger, and highly identifiable industrial actors (Cimprich et al., 2023). An original equipment manufacturer (OEM) occupying a central position in a B2B value chain is uniquely positioned to internalise a significant portion of this coordination. Rather than leaving material loops unmanaged across a fragmented network, the OEM can leverage its central position to orchestrate the ecosystem. Furthermore, because the OEM typically commands greater capital and operates under a longer investment horizon than smaller upstream or downstream partners, it is the actor best equipped to absorb the financial risks and high upfront costs identified by Hina et al. (2022).

This thesis adopts an OEM-centred decision-making perspective from Chapter 3 onward. Instead of modelling circular adoption as a diffuse, system-level phenomenon, the OEM is treated as the focal strategic actor whose business model choices directly shape material recovery outcomes, while still contextualising these choices within the broader stakeholder network described in Section 2.2.5.

2.3.3 Market Uncertainty and Commodity-Price Volatility in Circular Strategies

Unlike relatively stable determinants such as regulatory clarity or consumer awareness, commodity-price volatility is highly dynamic. Its financial impacts are substantial enough to warrant treatment as a distinct boundary condition shaping circular business model (CBM) adoption, rather than being aggregated within generic market barriers. Price fluctuations in raw materials directly undermine corporate cash flows and economic viability. Under these volatile conditions, conventional purchasing decisions based on static price comparisons can be highly misleading because they ignore market-risk exposure (Gaudenzi et al., 2021).

To quantify this volatility, Gaudenzi et al. (2021) propose integrating Total Cost of Ownership (TCO) with Real Options Valuation (ROV). Within this framework, TCO categorises costs into pre-transaction, transaction, and post-transaction phases, thereby capturing hidden operational expenses that a raw purchase price omits. Meanwhile, ROV evaluates the financial value of managerial flexibility, such as the option to switch to recycled materials when virgin commodity prices spike, under uncertainty that traditional

Net Present Value (NPV) metrics cannot address. This combined approach shows why a circular strategy's viability should not be assessed against a single, deterministic price projection, as this discards the strategic value of flexibility created by market volatility itself. The DSS developed in this thesis does not implement TCO or ROV, which require cost and cash-flow data that are rarely available in comparable form across alternative CBMs; it retains the underlying principle by re-evaluating the alternatives under explicit CRM market scenarios (Section 3.4).

This financial perspective aligns with broader critiques regarding how circular strategies interact with conventional market structures. Siderius and Zink (2023) argue that the circular economy represents a complex system of markets, and that circular initiatives underperform when constrained by traditional market characteristics, such as competition and private property. These forces naturally incentivise planned obsolescence and manufactured scarcity over technical efficiency.

Product-as-a-service (PaaS) configurations, specifically the use- and result-oriented archetypes discussed in Section 2.2.3, are not automatically transformative simply because the manufacturer retains ownership (Siderius & Zink, 2023). They remain subject to the same competitive pressures as linear models, achieving environmental and economic alignment only under specific market conditions. This qualification serves as an important corrective to the assumption that service-oriented business models represent an inherently safer economic path under volatile conditions. This caution is entirely consistent with the review by Tukker (2015), which confirms that the sustainability benefits of use- and result-oriented models are not empirically guaranteed.

Von Rennenberg et al. (2024) demonstrate how such market and policy uncertainties can be modelled explicitly rather than assumed away. Linking an Enterprise Input-Output (EIO) approach to the European Commission's criticality matrix, the authors evaluate the potential of urban mining to reinforce niobium circularity in Europe under three distinct future pathways: a resource-based supply bottleneck, a technology-driven sorting and recovery upgrade, and a government-mandated recycled-content quota.

Their analysis reveals that the resulting supply-risk reduction varies substantially across these three conditions, even though the nominal circular strategy remains identical (von Rennenberg et al., 2024). This structure of testing qualitatively different market and policy futures against a consistent set of circular strategies establishes a direct methodological precedent for the scenario-based sensitivity approach adopted in Chapter 3 and applied to NdFeB magnets in Chapter 5.

Methodologically, this thesis scales this logic from a resource-flow inquiry (quantifying overall supply-risk reduction) to a strategic business-model question: which CBM alternative remains economically preferable under varying levels of CRM price volatility. Taken together, Sections 2.3.1 to 2.3.3 provide the theoretical and empirical foundation for incorporating market uncertainty as an explicit, dynamic input to CBM evaluation, allowing the relative economic attractiveness of alternative CBMs to be assessed across different volatility conditions.

2.4 Decision Support Systems and Multi-Criteria Decision-Making in the Circular Economy

2.4.1 Foundations: What a DSS Is

In its foundational formulation, a Decision Support System (DSS) can be characterised as an interactive, model-based system designed to assist managers faced with semi-structured choices, decision-making environments where problems are specific enough to be partially formalised, yet too complex to be resolved by rigid, automated algorithms alone. Within the literature on circularity and reverse logistics, the development of such systems is not a recent phenomenon. Early quantitative reviews show that pioneering DSS frameworks, such as the REVLOG model in the late 1990s, were already being deployed to compare circular supply chain designs, assist in recovery network configuration, and evaluate “design-for-disassembly” strategies from an economic standpoint (Fleischmann et al., 1997). As the field transitioned towards Closed-Loop Supply Chains (CLSC), systematic reviews highlighted that decision-making under uncertainty and multi-criteria performance evaluation represent some of the most critical and highly researched domains in corporate environmental management (Govindan et al., 2015).

Modern applications have sought to formalise this decision support to address the qualitative vagueness and multi-criteria trade-offs inherent in circular strategies. A prominent framework is developed by Ahmed et al. (2022), who integrate fuzzy logic with multi-criteria decision-making (MCDM) methods. In their architecture, a fuzzy-logic layer translates subjective, linguistic judgements from managers into precise, crisp numbers, which are subsequently aggregated through techniques such as the Analytic Hierarchy Process (AHP) and TOPSIS into rankable circularity scores (Ahmed et al., 2022). This design bridges unstructured human intuition and complete automation, ensuring that expert qualitative knowledge is preserved rather than eliminated (Ahmed et al., 2022).

This tiered architecture, separating the data input layer, the modelling procedure, and the final decision output, is further supported by the evolution of digital-enabled circular frameworks. Indeed, the “Smart CE” model demonstrates that digital data pipelines, rooted in the classic Data-Information-Knowledge-Wisdom (DIKW) hierarchy, exist precisely to support, augment, and eventually unburden human decision-makers faced with systemic resource-efficiency challenges (Kristoffersen et al., 2020). In contemporary B2B settings, this has led to the emergence of AI-based DSS and expert systems (Deveci et al., 2024). These advanced decision frameworks leverage perceptive, predictive, and prescriptive analytics to help industrial manufacturers and OEMs navigate complex operational and strategic scenarios, such as optimising resource allocation or evaluating life-extension strategies like predictive maintenance (Sjödén et al., 2023). Ultimately, this layered and multi-criteria logic directly informs the decision-support framework proposed and formalised in Chapter 4 of this thesis.

2.4.2 Circularity Indicators and Assessment Frameworks, and Their Limits for Comparing CBM Alternatives

A substantial body of circular economy (CE) decision-support research has focused on measuring and assessing circularity, while comparatively less attention has been devoted to the direct comparison of alternative circular business model configurations. In their

systematic review of 55 sets of circularity indicators, Saidani et al. (2019) develop a ten-category taxonomy and link it to an Excel-based query tool designed to support decision-makers in navigating a highly fragmented methodological landscape. This fragmentation complicates the identification and selection of appropriate indicators for practitioners and highlights the need for more structured approaches to support the application of circularity assessment in industry (Saidani et al., 2019).

This pattern is further illustrated by Kristensen and Mosgaard (2020) in their review of 30 micro-level circularity indicators. They find that the reviewed indicators are predominantly oriented towards recycling and other end-of-life strategies, while inner-loop strategies such as reuse, repair, and maintenance receive comparatively limited attention. Furthermore, the majority of indicators focus on economic aspects, while environmental and particularly social dimensions receive less attention, limiting the extent to which existing metrics provide a balanced assessment across sustainability dimensions (Kristensen & Mosgaard, 2020). Even within this micro-level literature, limited attention is given to product functionality and sharing, potentially constraining the applicability of existing indicators to Product-Service Systems (PSS) and other service-oriented circular business models (Pollard et al., 2021).

Integrative frameworks, such as the multi-level assessment model developed by Ahmed et al. (2022), attempt to address this fragmentation by combining micro-, meso-, and macro-level perspectives within a unified fuzzy-logic and multi-criteria architecture. However, their primary focus remains on assessing and ranking the circularity of products or processes, rather than on selecting among alternative circular business model configurations for a specific recovery context. The distinction is important because the decision addressed in this thesis is not simply whether, or to what extent, a given configuration is circular, but which alternative business model is comparatively preferable under specified technological, economic, and market conditions.

A complementary perspective emerges from the corporate raw-material criticality literature. In their systematic review of corporate-oriented criticality assessments, Otterbach and Fröhling (2025) show that existing approaches provide valuable support for identifying material-related risks but offer more limited support for operationalising subsequent risk-mitigation decisions. Specifically, less than 5% of the environmental, social, and supply-risk indicators examined are company-specific, with many relying instead on generic or global-average data. The review also identifies a weak linkage between risk assessment and the subsequent selection of risk-mitigation strategies, with only limited consideration of circular-economy strategies beyond recycling (Otterbach & Fröhling, 2025). Taken together, these strands of literature point to a comparatively underdeveloped link between CRM risk assessment and the ex-ante comparative selection of circular business models under changing economic and market conditions. This is the specific gap addressed by this thesis through the development of a scenario-based DSS designed to compare alternative CBM configurations and assess their relative economic attractiveness under different CRM market conditions.

2.4.3 DSS Applications in the Circular Economy

Beyond circularity measurement narrowly defined, decision support for the circular economy (CE) has developed along several related, contemporary paths. A prominent, recent strand of literature frames decision support around artificial intelligence (AI) and digital technologies. Although this thesis does not employ an AI-based system, this digitally enabled

domain represents a highly relevant area of contemporary research. AI-based decision support systems are conceptualised as tools to identify optimal sustainable pathways and anticipate future trends under uncertainty (Deveci et al., 2024). Within this literature, digital technologies are consistently framed as critical enablers that translate circular intentions into implementable strategies rather than as substitutes for managerial discretion. This distinction is particularly relevant given empirical evidence that while approximately 70% of supply chain leaders plan to invest in circularity, only 12% have formally linked their digital and circular strategies (Kristoffersen et al., 2020). However, despite the advanced nature of this AI-oriented body of work, it exhibits a similar limitation to the one identified in Section 2.4.2. A bibliometric review of 821 publications at the intersection of AI and CE reveals that research remains heavily clustered around operational themes such as waste management, reverse logistics, and remanufacturing, leaving the critical questions of economic feasibility and strategic business model selection under market uncertainty largely unaddressed (Raut et al., 2025).

Two further streams of research, closer to industrial application, reinforce this point from complementary directions. First, while AI capabilities are theoretically positioned as efficiency catalysts for circular business model innovation (CBMI), a systematic review by Madanaguli et al. (2024) reveals that firms often lack the complex organisational routines and processes required to operationalise this potential. This organisational barrier suggests a clear role for more transparent, lightweight, and non-AI decision-support tools that firms can readily implement. Furthermore, in business-to-business (B2B) digital servitisation contexts Sjödin et al. (2023) distinguish between AI applications that automate decision-making and those that augment it (encompassing perceptive, predictive, and prescriptive support). They argue that successful CBMI in industrial settings depends on this augmentation logic specifically, whereby tools are designed to inform and enhance managerial choice rather than replace human agency (Sjödin et al., 2023). This augmentation philosophy directly aligns with the role of the scenario-based DSS developed in this study. Because circular decision-making at scale is inherently a systemic group process involving multiple stakeholders rather than a single-firm exercise (Modgil et al., 2021), a separate, participatory strand of literature emphasises that the variables and criteria feeding circular decision support frameworks should be co-defined with, and validated by, the relevant value chain actors (Pollard et al., 2021; Pollard et al., 2023). This stakeholder-centric validation ensures alignment with practical industry needs and supports the methodological choice made in this thesis, from Chapter 3 onwards, to refine and weight the proposed DSS through expert consultation.

2.4.4 Multi-Criteria Decision-Making and the Analytic Hierarchy Process

The circularity assessment applications reviewed in Sections 2.4.2 and 2.4.3 primarily support the measurement and assessment of circularity, rather than the strategic comparison of alternative circular business models (CBMs). Comparing CBMs for critical raw material recovery requires decision-makers to simultaneously weigh financial returns against reverse-logistics costs, regulatory exposure, value chain coordination, and material availability. These criteria are neither directly comparable nor reducible to one another without a structured multi-criteria decision-making (MCDM) methodology. The necessity of adopting MCDM approaches to manage these complex, multi-dimensional trade-offs under market and operational uncertainty is widely documented in comprehensive closed-loop supply chain reviews (Govindan et al., 2015).

The Analytic Hierarchy Process (AHP), originally developed by Saaty (1980) and subsequently refined for practical managerial application (Saaty, 1990), is designed precisely to address this class of semi-structured problems. It structures a complex decision into a hierarchical model consisting of three primary layers: the overall objective, the evaluation criteria, and the alternatives being compared. Through a series of pairwise comparisons made on a standardised nine-point scale, the method derives mathematical weights for each criterion and prioritised scores for each alternative (Saaty, 1990). The framework incorporates a built-in Consistency Ratio (CR) to verify that the pairwise judgements supplied by experts are internally coherent, automatically flagging inconsistent inputs that exceed the acceptable 10% threshold (Saaty, 1990).

Within the wider circular economy and recovery literature, systematic reviews document that AHP has been extensively applied to validate OEM decision frameworks, select sustainable end-of-life product recovery options, and configure remanufacturing technology portfolios (Govindan & Soleimani, 2017). Furthermore, because circular decision-making frequently relies on subjective qualitative inputs, AHP is often combined with complementary ranking methods to translate stakeholder-derived judgements into rankable performance scores across competing circular alternatives (Ahmed et al., 2022). This established integration confirms that the core mechanics of the method are highly transferable and robust when applied to circularity trade-offs.

2.5 Research Gap and Positioning of the Study

The four bodies of literature reviewed in this chapter have largely developed in parallel. Critical Raw Material risk and recovery, Circular Business Model typologies and their control and technical-recovery dimensions, the empirical determinants of CBM adoption together with commodity-price volatility, and decision-support and multi-criteria methods each offer valuable insight on their own terms. A firm evaluating circular options for CRM recovery, however, has to weigh all of these issues at once: price fluctuations, uncertain secondary-material availability, regulatory change, reverse-logistics cost, ownership structure, and investment feasibility, within a single decision, a task for which the four strands reviewed above, taken individually, offer only partial and indirect support.

Two limitations, which have been anticipated at several points across this chapter, are particularly relevant for positioning the contribution of this study. First, circular economy (CE) strategies are still often assessed without fully incorporating dynamic market conditions. As discussed in Section 2.3.3, commodity-price volatility can have measurable financial effects across supply chains, affecting pricing decisions, budget planning, net cash flows, and the value of risk-mitigation strategies, dimensions that static evaluations may not adequately capture (Gaudenzi et al., 2021). Circular initiatives remain embedded in wider market systems in which systemic cost-efficiency pressures, competition, and weak reverse-logistics incentives can constrain circular outcomes independently of a strategy's technical potential (Siderius & Zink, 2023). Second, as shown in Section 2.4.2, the literature provides a fragmented landscape of circularity indicators and assessment frameworks operating at different levels of analysis (Saidani et al., 2019; Kristensen & Mosgaard, 2020; Ahmed et al., 2022). These approaches are primarily concerned with assessing, measuring, or monitoring circularity performance and are therefore less directly oriented towards the earlier, ex-ante managerial decision of comparing alternative Circular Business Model (CBM) configurations and determining their relative suitability for a specific firm and operating context.

The research gap emerging from Sections 2.1 to 2.4 does not primarily concern a lack of technical recycling solutions or CBM typologies, both of which are comparatively well established in the literature. Existing approaches remain fragmented across the dimensions relevant to the strategic selection of circular business models. The literature provides limited integration between the comparative evaluation of alternative CBM configurations, the perspective and decision requirements of a specific industrial actor, the multiple economic, market, organisational, technical, and regulatory factors that shape CBM feasibility, and changing market conditions affecting the availability and economic value of critical raw materials. A structured Decision Support System (DSS) can therefore provide a means of bringing these dimensions together within an ex-ante decision process, supporting the comparison of alternative CBM configurations under explicit market scenarios rather than under a single, fixed set of assumptions.

This need is further supported by research on criticality assessment and circular business models. Recent work on raw-material criticality highlights the importance of connecting criticality considerations with decisions concerning circular strategies, while indicating that existing assessment approaches remain primarily focused on material criticality and recycling-related dimensions rather than on the comparative evaluation of alternative business-model configurations (Otterbach & Fröhling, 2025). Similarly, research on circular and digitally enabled business models has examined the development and implementation of circular and servitisation-oriented configurations in manufacturing contexts, but these contributions provide a different analytical focus from the CRM-intensive, firm-level strategic selection considered in this thesis (Madanaguli et al., 2024; Sjödin et al., 2023). Taken together, these streams reinforce the observation that the relevant elements of the decision problem are addressed across different, largely separate bodies of literature.

Positioning this thesis relative to the identified gap requires specifying both the focal decision-maker and the strategic alternatives considered. Section 2.2.6 and Section 2.3.2 established the relevance of an original equipment manufacturer (OEM) occupying a central position in a business-to-business value chain, such as an offshore wind turbine manufacturer, particularly because of its influence over product architecture, customer relationships, recovery arrangements, and the investments required to implement circular business models (Cimprich et al., 2023). This positions the OEM as a particularly relevant actor for the integration outlined above: decisions concerning circular business models in CRM-intensive industries are closely connected to the strategic management of material supply risk, recovery opportunities, investment requirements, and exposure to CRM market conditions. The OEM is therefore considered in this thesis as the focal decision-maker responsible for evaluating alternative CBM configurations in a CRM-intensive industrial context. Section 2.2.3 defined the three CBM configurations considered in this study, product-, use-, and result-oriented Product-Service Systems (Tukker, 2004). These configurations constitute the strategic alternatives evaluated by the focal OEM.

To integrate these dimensions, Chapters 3 to 5 of this thesis develop a scenario-based, managerially oriented Decision Support System (DSS) for OEMs operating in CRM-intensive industries. The framework evaluates and compares the three CBM configurations using the factor set and trade-off logic derived from the actor-mapping approach discussed in Section 2.2.5 (Konietzko et al., 2020a). The alternatives are evaluated under different market-uncertainty conditions introduced in Section 2.3.3 (Gaudenzi et al., 2021; Siderius

& Zink, 2023), while the multiple evaluation criteria are aggregated through the Analytic Hierarchy Process presented in Section 2.4.4 (Saaty, 1990).

The contribution of the thesis lies in integrating established knowledge on CRM circularity, CBM configurations, adoption factors, market uncertainty, and multi-criteria decision-making into a single, structured DSS framework. This framework supports an OEM in a CRM-intensive industry in the ex-ante comparison and selection of alternative CBM configurations under changing market conditions. Rather than replacing existing CBM typologies, assessment methods, or MCDM techniques, the framework connects these elements to a specific managerial decision and provides a structured basis for evaluating how the preferred CBM may change across alternative CRM market scenarios.

3. Methodology

The research problem addressed in this thesis is constructive rather than explanatory: the objective is not to test why OEMs behave in a given way when circular business models are considered, but to develop a decision-support artefact capable of guiding an OEM's choice among Product-, Use- and Result-oriented CBMs under conditions of CRM scarcity and volatile market conditions. This choice is inherently multi-dimensional, since the relative suitability of the three archetypes depends simultaneously on economic, technical, organisational, market-related and regulatory considerations that cannot be reduced to a single financial or operational indicator. A structured methodological approach is therefore required, one capable of integrating heterogeneous decision criteria, of eliciting the relative importance that experts attach to them, and of testing how the resulting recommendation behaves once external conditions are allowed to vary.

The remainder of this chapter sets out that approach in five stages, moving from the paradigm within which the research is positioned (Section 3.1), to the derivation of the decision variables (Section 3.2), to the analytical procedure through which those variables are weighted (Section 3.3), to test the sensitivity of the resulting recommendation to CRM market scenarios (Section 3.4), and closes with the empirical collection of expert judgements used to review the framework and derive the AHP weights (Section 3.5).

3.1 Overall Research Design: A Design Science Research Approach

An explanatory or interpretive research design would not, on its own, directly address the primary objective of this thesis, which is to develop a decision-support artefact for an OEM's choice among alternative CBM archetypes. The overall research design is therefore framed as a Design Science Research (DSR) project, a paradigm explicitly concerned with the construction and rigorous evaluation of artefacts that address a problem of practical relevance while remaining grounded in an existing body of theoretical and methodological knowledge (Hevner et al., 2004). Where the behavioural-science tradition that dominates much of management research seeks to develop and justify theories that explain or predict organisational phenomena, design science seeks to develop and justify artefacts, understood broadly as constructs, models, methods and instantiations, that extend what an organisation or an individual is able to do (Hevner et al., 2004).

The artefact developed in this thesis is a scenario-based Decision Support System (DSS) for the selection of the most suitable CBM archetype by an OEM operating under volatile CRM market conditions. The three CBM archetypes examined, the consolidated set of decision variables built around them, the six-category hierarchy into which those variables are organised, and the resulting Analytic Hierarchy Process (AHP) architecture together constitute the DSS: primarily a model, in the sense of a set of constructs and their relationships intended to represent the OEM's decision problem, coupled with a method through which that model is populated with expert judgement, weighted, and used to compare the three archetypes for a given decision context. The single case study developed in Chapter 5 provides an initial instantiation of the DSS, demonstrating how the framework can be operationalised in a specific CRM-intensive application context rather than remaining a conceptual proposal (Hevner et al., 2004; Gregor & Hevner, 2013).

Hevner (2007) situates the design process at the intersection of three interconnected cycles. The relevance cycle connects the design activity to the application environment, supplying the requirements the DSS must satisfy and the criteria against which it will later be evaluated. In this thesis that cycle is fed by the managerial and industrial problem set out in Chapter 1 and by the literature on CRM criticality, CBM adoption barriers and market volatility reviewed in Chapter 2, which together establish why OEMs need a structured way of comparing CBM archetype alternatives and what such a comparison must be able to capture. The rigor cycle, in Hevner's (2007) framework, connects the design activity to the existing knowledge base, ensuring that the DSS is grounded in established theoretical and methodological knowledge. The knowledge base is grounded chiefly in Tukker's (2004) typology of product-, use-, and result-oriented product-service systems, developed in full as the three CBM archetypes in Section 4.1, complemented by the broader CBM literature and by research on servitisation and product-service systems concerning ownership, risk, responsibility and revenue allocation (Mont, 2002; Baines et al., 2007; Reim et al., 2015). The design cycle concerns the iterative construction and refinement of the DSS and is documented across Chapters 3 to 5: an initial factor set is derived from the structural comparison of the archetypes and consolidated against the literature, reorganised into a criteria hierarchy, weighted through expert elicitation, and finally demonstrated and examined through a single illustrative case.

This cyclical process aligns with the six-activity Design Science Research Methodology proposed by Peffers et al. (2007): problem identification and motivation, definition of the objectives for a solution, design and development, demonstration, evaluation, and communication. These six activities provide a process logic for structuring the research described in this thesis. Problem identification and motivation, together with the definition of the objectives pursued through the three research questions, are addressed in Chapters 1 and 2. Design and development, understood as the construction of the decision variables, the categorical hierarchy, the weighting procedure and the scenario mechanism, is undertaken across Chapters 3 and 4. Demonstration is undertaken in Chapter 5 through the application of the DSS to a single illustrative case. Evaluation is addressed through an expert review of the completeness, relevance and clarity of the decision-variable structure, informed by the literature on which the artefact is based, and through the examination of DSS behaviour under alternative CRM market scenarios. Communication is addressed through the discussion (Section 5.6) and conclusions in Chapter 6.




Fig. 3.1 - DSRM Process Model (Peffers et al., 2007)

The choice of Design Science Research over a qualitative case-study design or a quantitative survey-based design reflects the nature of the intended contribution. Following the knowledge contribution framework proposed by Gregor and Hevner (2013), this thesis is positioned as an improvement-type contribution. It does not claim to introduce an entirely new class of solution, since CBMs, the Analytic Hierarchy Process and scenario-based sensitivity analysis are each individually well established, but it proposes a new configuration of these known techniques, an OEM-centred, archetype-driven decision-variable set combined with a hierarchical weighting procedure and a scenario mechanism for CRM market conditions, applied to a decision problem for which no comparable framework was identified in the literature reviewed in Chapter 2. Extending Design Science Research beyond its original information-systems setting into management and operations research contexts, Dresch et al. (2015) note that this kind of improvement contribution is an entry point for design-oriented management research, where the objective is less to produce a generalisable scientific law than to produce an artefact that addresses a specific, well-bounded class of organisational problems more effectively than the alternatives currently available. That is the standard against which the framework developed in this thesis should be read.

Consistent with the Design Science Research paradigm adopted in this section, the artefact is not evaluated through statistical generalisation from a representative sample. Instead, its operation and applicability are examined through a demonstration in a purposively selected case, complemented by the evaluation activities described in Sections 3.5 and Chapter 5 (Hevner et al., 2004). Chapter 5 sets out, at its opening, the rationale for adopting a single, illustrative case study as the demonstration context for the DSS, together with the criteria used to select that case. The present chapter is limited to establishing that such an evaluation approach is the one consistent with the overall research design just described.

A final remark on the scope of the resulting DSS is warranted before turning to its construction. As is typical of improvement-type design science contributions built around a single illustrative case, what this thesis offers as broadly transferable is the architecture of the DSS itself. The elements intended to be transferable are the logic used to derive decision variables from a structural comparison of the CBM archetypes, the six-category hierarchy into which those variables are organised, the AHP-based weighting and scoring procedure, and the mechanism used to test scenario sensitivity. What remains case-specific is the numerical content populating that architecture, namely the particular set of expert judgements collected and the resulting variable weights. Applying the framework to a different CRM-intensive asset would be expected to reuse the same decision variables and hierarchy while requiring a fresh elicitation of weights appropriate to that context.

3.2 Method for Factor Identification

This section sets out how the decision factors and decision variables that populate the DSS's criteria structure were identified from the literature. The section proceeds in six steps. First, it fixes the analytical lens through which the three CBM archetypes are examined and identifies the actor whose interests define what counts as a relevant factor (Section 3.2.1). Second, it maps, for each archetype separately, the value chain described in the literature across the product life cycle (Section 3.2.2). Third, it compares the three value chains to identify factors generated by their structural differences and separately identifies contextual conditions that do not vary across archetypes but must nevertheless be considered when assessing how each archetype performs within a given decision context (Section 3.2.3).

Fourth, on the basis of Assmann et al. (2023) and Geissdoerfer et al. (2023), it defines six macro-categories for organising the resulting factors (Section 3.2.4). Fifth, it consolidates the candidate factors into the final set of decision variables that forms the criteria structure of the MCDA engine (Section 3.2.5). Sixth, it explains how each decision variable is rated across the three CBM archetypes and how its benefit/cost classification is recorded for the numerical procedure developed in Section 3.3.4 (Section 3.2.6).

3.2.1 The OEM as Focal Actor and the Common Analytical Lens across the Three Archetypes

The starting point of the analysis is Tukker's (2004) typology of product-service systems (PSS), which provides the basis for the three CBM archetypes examined in this thesis: product-oriented, use-oriented and result-oriented. Throughout this thesis, "CBM" denotes the general concept of a circular business model, while "CBM archetype" refers specifically to one of these three configurations. Once introduced the two terms are used interchangeably wherever the referent is unambiguous from context. The three differ in the degree of servitisation involved, in whether ownership of the physical asset is transferred to the customer or retained by the provider, and in how far responsibility for the asset's performance and end-of-life fate extends beyond the point of sale (Tukker, 2004). These three archetypes were selected because they represent qualitatively distinct positions along the same underlying continuum, and because that continuum is precisely what determines how risk, revenue and material responsibility are distributed between provider and customer once ownership is not transferred outright (Mont, 2002; Baines et al., 2007; Reim et al., 2015).

Because the three archetypes redistribute risk and responsibility differently, their comparison must be conducted from the perspective of a specific actor whose interests define the decision problem. This thesis therefore adopts the OEM as the focal actor, consistent with the role attributed in the critical-raw-material literature to manufacturing firms that retain influence over the product value chain (Cimprich et al., 2023). Framing the OEM as the focal actor determines which factors are considered relevant: a factor is retained only if it bears on the feasibility, attractiveness or implementation of a CBM from the standpoint of the actor that must design and finance the offer and, under the use- and result-oriented archetypes, operate or assume responsibility for the asset, rather than on circularity outcomes in general or on the experience of the end customer. Adopting a single focal actor and a fixed set of three alternatives also allows the CBM archetypes to be examined through a common analytical lens rather than through three independently developed lists of factors. Throughout Sections 3.2.2 and 3.2.3, the guiding question is which conditions make an OEM more or less able, and more or less willing, to adopt each archetype relative to the other two. Establishing this common lens before the value chains are mapped ensures that the three archetypes are analysed on directly comparable terms.

3.2.2 Value Chain Mapping of Each CBM Archetype across the Product Life Cycle

Before the three archetypes can be compared, each is mapped on its own terms as a self-contained value chain, identifying the actors involved, the activities each of them performs, and the degree of control the OEM retains or relinquishes at each step. This sequencing matters because a comparison attempted before the individual value chains are established would risk comparing activities that are only nominally the same. Maintenance, for instance, is carried out under a materially different governance arrangement, and with a different

allocation of cost and liability, depending on whether the asset has been sold, leased or supplied as a guaranteed outcome (Reim et al., 2015; Baines et al., 2007). Mapping each archetype separately first therefore fixes a common vocabulary of actors and activities against which genuine points of divergence, rather than superficial differences of description, can later be identified.

Each value chain is unpacked along three life-cycle macro-phases, beginning-of-life (BOL), middle-of-life (MOL) and end-of-life (EOL), a partition consistent with the way the circular economy literature on critical raw materials frames the strategic decisions available to a manufacturer across a product's life. Design and sourcing choices made before the asset ever reaches the customer, use and maintenance decisions that shape how intensively and for how long it is exploited, and collection, recovery and residual-disposal decisions that determine whether the materials it contains re-enter the economy at all (Gaustad et al., 2018; Bocken et al., 2016). Within each phase, the mapping records, activity by activity, who acts, what interest or responsibility the OEM has in that activity, and whether the OEM controls it directly or depends on another actor to carry it out. The resulting three accounts, one per archetype, are deliberately left uncombined at this stage, consolidating and contrasting them is the task taken up in Section 3.2.3.

3.2.3 Cross-Model Comparison: Structural and Context-Derived Decision Factors

With the three value chains mapped on comparable terms, the archetypes are now set directly against one another, phase by phase, to isolate the points at which they genuinely diverge rather than simply differ in description. Each such divergence is translated into a candidate decision factor: a condition that changes systematically according to which of the three CBMs is chosen and that an OEM would therefore need to weigh when comparing them. Because ownership of the asset transfers to the customer under the product-oriented model but is retained by the provider under the use- and result-oriented ones, for example, capital intensity, the timing of revenue collection, and the degree of actual control over whether the asset is eventually returned all vary directly with the archetype selected (Tukker, 2004; Reim et al., 2015; Gaustad et al., 2018). Factors of this kind are labelled structural because their value is a direct consequence of the CBM chosen.

Relying on the cross-model comparison alone, however, would leave the resulting factor set structurally incomplete, since a procedure built to isolate divergence can, by construction, only surface conditions that differ across the three archetypes. Several conditions drawn from the same literature matter just as much to the choice among the three CBMs without themselves varying across them. The prevailing volatility of the critical raw material market, the willingness of customers to accept non-ownership consumption, the maturity of take-back regulation, and the reverse-logistics and processing capacity available in the wider industrial ecosystem are all fixed, at any given moment, regardless of which CBM the OEM selects (Cimprich et al., 2023; Gaustad et al., 2018). Precisely because the cross-model comparison cannot surface them on its own, these conditions were identified separately, directly from the literature, and deliberately incorporated into the same decision-factor set since assessing how each of the three archetypes performs once such conditions are held constant is no less important to the comparison than assessing the archetypes' intrinsic structural differences. For example, an OEM choosing among the three CBM archetypes needs to know not only how the three differ by design, but how each of them is expected to behave once demand is uncertain, once the raw material market is volatile, or once take-back

regulation is weak. Factors of this second kind are labelled context-derived, since they describe a state of the environment rather than a consequence of the decision itself, but they are decision factors in the same sense as the structural ones.

The three CBMs are compared against them just as they are compared against the structural factors, only their value does not itself change with the archetype chosen. The distinction mirrors one already established in the strategic decision-making literature, between elements that lie within the scope of a given choice and predetermined elements external to it, whose value must instead be treated as an assumption, or later as a scenario, against which the choice is evaluated (Schoemaker, 1995).

This comparison therefore distinguishes three, rather than two, analytical roles that a candidate factor can occupy. Neither the structural nor the context-derived role, on its own, explains why a condition such as CRM price volatility can later reappear, in Section 3.2.5, as a criterion the three archetypes are scored against: the condition itself does not vary with the archetype chosen, so it cannot by itself differentiate Product-, Use- and Result-oriented CBM archetypes. What differentiates them is the archetype-specific exposure, burden or benefit that the same contextual condition produces once it is read through each archetype's own structural characteristics. This archetype-specific exposure is the third analytical role, a context-derived decision variable, and it is what Section 3.2.5 shows how to construct out of a contextual condition that does not itself vary by archetype.

Tagging every candidate factor as structural or context at the point it is first identified, rather than only when scenario analysis is introduced later in this thesis, still has a direct methodological payoff: it keeps explicit, for every decision variable, why its baseline rating differs across the three archetypes in the first place, whether that is a direct consequence of the archetype's own configuration or an archetype-specific reading of an external condition. The CRM market scenarios developed in Section 3.4 are modelled as a change to whichever decision variables can be shown to respond causally to that condition, a property established independently of this structural/context-derived tagging. The CBM performance assessments attached to scenario-sensitive variables are what a scenario is allowed to move, while the importance weights elicited from experts through the AHP remain fixed across scenarios. The cross-model comparison conducted in this section therefore produces not a single undifferentiated list of factors, but two tagged families of candidates that are carried forward, together, into the categorisation and consolidation steps described next.

3.2.4 A Literature-Derived Coding Grid: Deriving Six Macro-Categories

Before the candidate factors generated in Section 3.2.3 can be compared through the AHP, they need to be organised into a small number of macro-categories, both to keep the decision problem manageable for an expert evaluating it and to ensure that the resulting structure maps consistently onto the wider literature on CBM adoption. Six macro-categories were defined on this basis.

Assmann et al. (2023) analyse fifty-four determinants extracted from sixty-seven journal articles and organise them into eight macro-categories: culture, regulation, market, strategy, business case, collaboration, operations and knowledge. Geissdoerfer et al. (2023) analyse a comparable set of drivers and barriers, drawn from a structured review combined with a multi-case study of twenty-one organisations, and organise them into seven categories that include market, financial, technical, organisational, value chain and legal determinants, among others. Based on these two reviews, six macro-categories were defined for the

purposes of this thesis. Both reviews follow the kind of transparent, protocol-driven review process recommended for building evidence-informed classification schemes in management research (Tranfield et al., 2003), which is what makes them a defensible basis for the resulting coding grid.

Because the two reviews were not built to be mutually consistent with one another, comparing them directly required an explicit act of interpretation: some pairs of categories addressed closely related concerns from different angles and were merged into a single macro-category, such as Assmann et al.'s (2023) business case category and Geissdoerfer et al.'s (2023) financial category, both concerned at bottom with investment, cost and profitability. Other categories, such as Assmann et al.'s (2023) operations category, were broader than any single category on the other side of the comparison and were therefore distributed across more than one macro-category according to the actual content of their determinants, rather than kept whole. This kind of comparison and aggregation of independently produced classification schemes into a smaller number of higher-order categories follows the same logic that qualitative research methodology recommends when moving from first-order source material to second-order, researcher-constructed themes. The resulting categories must remain traceable back to the specific source content that generated them rather than being introduced as free-standing labels (Gioia et al., 2013).

The comparison yields six macro-categories, each capturing a distinct dimension of the OEM's decision problem: Market, grouping demand-side and competitive conditions; Economic/Business Case, grouping the CBM's financial architecture, from investment requirements and operating costs to value capture and exposure to CRM price volatility; Technical/Product, grouping the characteristics of the asset and of the technology needed to deliver and verify the offer; Organisational/Capabilities, grouping what the OEM itself must be able to do, as distinct from what the market or the asset determines; Value Chain/Ecosystem, grouping the OEM's relationships with actors outside its own boundary and the physical and informational flows connecting them; and Regulatory/Institutional, grouping the legal and policy environment within which the OEM operates.

Macro-category	Assmann et al. (2023)	Geissdoerfer et al. (2023)	Rationale for the mapping
Market	Market	Market	Retained as in both sources
Economic / Business case	Business case	Financial	Conceptual merger: both concern resources, investment, cost and profitability
Technical / Product	Operations (product-related determinants)	Technical	Narrowed to product- and production-related characteristics
Organisational / Capabilities	Culture, Strategy, Knowledge, Operations (organisational-resource determinants)	Organisational	Aggregation of factors internal to the firm
Value chain / Ecosystem	Collaboration, Operations (supply-chain-related determinants)	Value chain	Relationships, flows and actors external to the firm
Regulatory / Institutional	Regulation	Legal	Conceptual merger of external rule-setting institutions

Tab. 3.1 - The six macro-categories from Assmann et al. (2023) and Geissdoerfer et al. (2023).

Two coding rules keep the resulting grid from fragmenting once it is applied to the factor pool. First, each factor is assigned to a single primary macro-category. A genuine secondary connection to another category is recorded as a note rather than treated as a second, parallel assignment, since allowing one factor to sit under two categories at once would reintroduce, inside the very structure meant to remove it, the same inconsistency that motivated deriving the grid independently in the first place. Second, recovery-related conditions are not given a seventh category of their own, even though recovery is central to this thesis's concern with critical raw materials. Instead, recoverability is treated as an attribute distributed across the six categories according to its actual content: a technical characteristic under Technical/Product, a reverse-logistics arrangement under Value Chain/Ecosystem, a take-back obligation under Regulatory/Institutional, and so on. This follows directly from how the servitisation literature itself describes recovery. A provider paid for a service or for a result, rather than for the product sold outright, retains a direct financial interest in the components it will eventually reclaim, so the incentive to prolong service life and recover value at end-of-life runs through the whole of a CBM's value chain rather than sitting apart from it as an independent concern (Tukker, 2015).

3.2.5 Consolidation of Decision Factors into the Final Set of Decision Variables

The candidate factors generated by the cross-model comparison in Section 3.2.3 were assembled one phase, and one archetype divergence, at a time, which is precisely why they

cannot be carried into the grid described in Section 3.2.4, and from there into the AHP, unchanged. A decision factor at that earlier stage is the qualitative concept identified either through the structural comparison of the archetypes or through one of the contextual conditions considered relevant to the OEM's decision, as set out in Section 3.2.3. A decision variable is what that concept becomes once it has been checked against every other candidate for overlap or redundancy, assigned to the macro-category its content actually belongs to, and reframed as a single evaluative dimension an OEM can be expected to score consistently across all three CBMs. Because the performance assessment adopted in this thesis rates each archetype, variable by variable, against a small number of ordinal categories, Low, Medium and High, mapped once onto a common scale rather than measured through precise numerical inputs (Saaty, 2008), a decision variable does not need to survive as a narrow formula requiring data that are rarely available in comparable form across the three CBM archetypes; it needs instead to be a construct the literature can support with a clear ordinal judgement for each archetype. This shifts the consolidation task from checking whether a specific indicator is well specified to checking whether the underlying construct is genuinely distinct from its neighbours and free of the kind of overlap that would let the same underlying condition weigh on the outcome twice. Applying the third analytical role identified in Section 3.2.3 to a contextual condition yields an evaluative construct ratable Product by Product, Use by Use and Result by Result. This is how critical raw material price volatility becomes the decision variable CRM Price Volatility Exposure, and how customer reluctance toward non-ownership consumption becomes Customer Acceptance of the CBM's Ownership/Access Arrangement.

Reviewing the candidate factors against this standard, one macro-category at a time, produces two recurring kinds of correction. The first is a boundary correction, applied when a factor had been placed in a macro-category adjacent to the one its content belongs to. Technical/Product and Organisational/Capabilities are the pair most exposed to this kind of correction, since a factor describing which technologies a CBM configuration requires is a property of the product and its supporting infrastructure, while a factor describing whether the OEM can actually exploit that infrastructure is a property of the organisation, and the two are readily conflated unless kept explicitly apart. The second is a construct-level consolidation, applied when several factors recorded separately during the cross-model comparison turn out, on inspection, to be different ways of answering the same underlying question about the OEM's exposure or capability. Comparing them as though they were independent criteria would let a single underlying condition be counted more than once, so they are instead folded into one broader evaluative dimension that a rater can still assess archetype by archetype on the strength of the same body of evidence. This consolidation operationalises, rather than formally tests, the independence AHP hierarchies assume between the elements being compared (Saaty, 1980, 2008): no statistical or correlation-based check was performed on the final set of decision variables, and this limitation is discussed in Chapter 6.

Applying this standard consistently across all six macro-categories yields the final set of decision variables used throughout the remainder of this thesis. Each variable is assigned to exactly one macro-category, is defined so as to be ratable on the Low/Medium/High scale for each of the three archetypes on the strength of the literature reviewed and remains traceable back to the specific structural or contextual divergence identified in Section 3.2.3 that first generated it. The resulting six-category hierarchy, populated with this final set of

decision variables, constitutes the criteria structure of the multi-criteria decision analysis engine developed in the remainder of this chapter. The decision variables themselves, and the macro-categories into which they are organised, are derived from the literature on CBM adoption rather than from the unconstrained judgement of the experts later asked to evaluate them. Before being carried into the AHP of Section 3.3, this literature-based set was submitted to the interviewed experts for a confirmatory check (Section 3.5.1), no omissions or misspecifications were identified. What those experts contribute through the AHP itself is therefore not the identification of the criteria, which the confirmatory check left unchanged, but the relative importance attached to criteria that have already been fixed (Saaty, 2008).

3.2.6 The Decision Variable×CBM Matrix: Literature-Based Ratings and Benefit/Cost Classification

Once the final set of decision variables and the six-category hierarchy are fixed, each variable is rated against each of the three CBM archetypes on the same ordinal scale, Low, Medium and High, producing a decision variable×CBM matrix. One row per decision variable, grouped by macro-category, one column per archetype, and one qualitative rating per cell. Each rating is assigned on the strength of the structural logic established in the cross-model comparison of Section 3.2.3 and of the same literature already used to identify and consolidate the variable in question, following the archetype-specific evidence set out where the variable is defined in Section 4.3. Where that literature does not itself provide a direct ordinal ranking between archetypes, the rating represents an explicit comparative judgement derived from the documented archetype characteristics rather than an empirical measurement or the researcher's unconstrained intuition. Formally, denoting by i the decision variable and by j the CBM archetype under consideration, every cell of the matrix records a rating drawn from the same three-level ordinal domain, as expressed in equation (3.1).

$$L_{i,j} \in \{Low, Medium, High\} \quad (3.1)$$

where:

- i = index of the decision variable ($i = 1, \dots, n$);
- j = index of the CBM archetype ($j \in \{\text{Product-oriented, Use-oriented, Result-oriented}\}$);
- $L_{i,j}$ = qualitative rating assigned, on the strength of the literature, to variable i for archetype j .

This matrix is the raw material that is reworked, in Section 3.3.4, into the CBM performance matrix used by the multi-criteria decision analysis engine. The Low/Medium/High matrix constructed in this section represents the baseline comparative performance of the three archetypes. The scenario analysis developed in Section 3.4 subsequently reassesses only those performance ratings whose underlying mechanism is causally sensitive to the selected CRM market scenario; every other rating in the baseline matrix is retained unchanged.

At this stage, however, a Low, Medium or High rating records only the intensity the literature attributes to a given condition for a given archetype. It does not yet record whether that intensity is desirable or undesirable from the OEM's perspective, and the two do not coincide for every variable. For a variable such as capital intensity or economic and price risk exposure, a High rating describes a condition that works against the attractiveness of the corresponding CBM. For a variable such as revenue generation potential or product recoverability, a High rating describes a condition that works in its favour. Combining ratings of both kinds in a single additive scoring procedure without first accounting for this

difference would treat two opposite conditions as though they pointed in the same direction, a problem long recognised in the multi-criteria decision-making literature through the distinction between benefit-type and cost-type criteria, whose values must be normalised in opposite directions before they can be aggregated (Hwang & Yoon, 1981). The matrix constructed at this stage therefore records, for every decision variable, not only its Low/Medium/High rating for each archetype but also a benefit/cost classification, established once from the definition of the variable itself rather than from any particular rating it happens to receive for a given archetype, as expressed in equation (3.2).

$$\delta_i \in \{benefit, cost\} \quad (3.2)$$

where:

- δ_i = benefit/cost classification of decision variable i , indicating whether a higher rating is favourable (benefit) or unfavourable (cost) to the attractiveness of the CBM archetype it is being assessed against.

How this classification is used to convert the raw ratings onto a single, consistently oriented numerical scale is described in Section 3.3.4.

3.3 The MCDA Decision Engine

The six-category hierarchy of decision variables established in Section 3.2 fixes what the multi-criteria decision analysis (MCDA) engine compares. It does not say how much weight each category or each variable within it should carry, nor how the literature-based Low/Medium/High ratings recorded in the decision variable×CBM matrix of Section 3.2.6 are to be combined into a single score for each of the three archetypes. This section develops the procedure that closes that gap. It proceeds in five steps: it first sets out the two-level hierarchical structure through which the decision variables are compared and the rationale for adopting it rather than a single, undifferentiated comparison of all variables at once (Section 3.3.1). It then describes how pairwise comparisons are elicited from experts and converted into local and global weights (Section 3.3.2) and it sets out how the internal coherence of those judgements is checked, both for each expert individually and for the aggregated group judgement (Section 3.3.3). It shows how the qualitative ratings and benefit/cost classification of Section 3.2.6 are converted into a single, consistently oriented numerical performance matrix (Section 3.3.4) and it closes with the weighted scoring rule used to rank the three CBM archetypes for a given decision context (Section 3.3.5). The procedure is based on the Analytic Hierarchy Process (AHP), a method for structuring and weighting multi-attribute decision problems through pairwise comparison and eigenvector-based priority derivation (Saaty, 1980, 2008).

3.3.1 AHP Hierarchy and the Rationale for a Two-Level Structure

The Analytic Hierarchy Process addresses a multi-attribute decision problem by decomposing it into a hierarchy of manageable elements, deriving priorities for the elements at each level of the hierarchy through pairwise comparison, and combining those priorities into an overall ranking of the alternatives under evaluation (Saaty, 1980, 2008). The hierarchy adopted in this thesis has two levels. At the first level, the six macro-categories defined in Section 3.2.4 are compared pairwise against one another, yielding a category weight for each. At the second level, and independently within each category, the decision variables belonging to that category are compared pairwise against one another, yielding a local weight for each variable relative to the other variables in the same category.

Adopting this two-level structure, rather than asking an expert to compare all decision variables directly in a single undifferentiated matrix, is not just a matter of convenience. The number of pairwise comparisons required to compare n elements directly grows with the square of n , as expressed in equation (3.3).

$$N_{\text{direct}} = \frac{n(n-1)}{2} \quad (3.3)$$

where:

- n = total number of decision variables in the hierarchy.

Under the two-level hierarchical structure adopted here, with six macro-categories and n_k decision variables in category k , the number of comparisons required is instead given by equation (3.4).

$$N_{\text{hierarchical}} = \frac{6(6-1)}{2} + \sum_{k=1}^6 \frac{n_k(n_k-1)}{2} = 15 + \sum_{k=1}^6 \frac{n_k(n_k-1)}{2} \quad (3.4)$$

where:

- n_k = number of decision variables belonging to macro-category k ($k = 1, \dots, 6$);
- the first term (15) is the number of pairwise comparisons among the six macro-categories themselves, at the first level of the hierarchy;
- the second term is the sum, category by category, of the pairwise comparisons required at the second level of the hierarchy.

Because n_k is by construction smaller than n for every category with more than one member, the hierarchical count of equation (3.4) is substantially smaller than the direct count of equation (3.3) whenever the decision variables are distributed across more than a single category, which is precisely the case for the hierarchy developed in this thesis. The cognitive burden of maintaining coherent judgement across a pairwise comparison matrix grows with its size, so a hierarchical decomposition of the kind recommended within the AHP literature itself is what keeps the elicitation task, and the internal consistency of its results, tractable (Saaty, 2008). Both expressions in equations (3.3) and (3.4) are kept in fully symbolic form, in n and n_k , since the final number of decision variables and their distribution across the six macro-categories are fixed only in Chapter 4, where the consolidation described in Section 3.2.5 is applied.

3.3.2 Pairwise Comparison and Weight Elicitation

At each level of the hierarchy, and separately for each macro-category at the second level, an expert compares the relevant elements two at a time, expressing for each pair the intensity with which one is judged more important than the other. These judgements populate a pairwise comparison matrix A , as expressed in equation (3.5).

$$A = \begin{bmatrix} 1 & a_{12} & \cdots & a_{1n} \\ 1/a_{12} & 1 & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ 1/a_{1n} & 1/a_{2n} & \cdots & 1 \end{bmatrix} \quad (3.5)$$

where:

- a_{ij} = intensity of preference expressed by the expert for element i over element j , measured on Saaty's fundamental 1-9 scale (1 = equal importance; 3 = moderate; 5

= strong; 7 = very strong; 9 = extreme; the even values 2, 4, 6 and 8 represent intermediate judgements) (Saaty, 1980, 2008).

The matrix is reciprocal by construction: an element is judged equally important as itself, and the judgement in one direction is the arithmetic reciprocal of the judgement in the opposite direction, as expressed in equation (3.6).

$$a_{ii} = 1, a_{ji} = \frac{1}{a_{ij}} \quad (3.6)$$

The vector of priorities associated with the matrix A, denoted w, is derived through the principal eigenvector method: w is the eigenvector of A associated with its principal (largest) eigenvalue $\lambda_{\max}$ normalised so that its elements sum to one, as expressed in equations (3.7) and (3.8) (Saaty, 1980, 2008).

$$A\mathbf{w} = \lambda_{\max}\mathbf{w}, \mathbf{w} = \begin{bmatrix} w_1 \\ w_2 \\ \vdots \\ w_n \end{bmatrix} \quad (3.7)$$

$$\sum_{i=1}^n w_i = 1 \quad (3.8)$$

where:

- w = column vector of priorities (local weights), of dimension equal to the number of elements compared in A;
- w_i = priority, or local weight, of element i;
- $\lambda_{\max}$ = principal eigenvalue of A, used again in the consistency check of Section 3.3.3.

This procedure is applied at both levels of the hierarchy: once to the matrix comparing the six macro-categories, yielding the category weights W_c , and once, independently, to each of the six matrices comparing the decision variables within a category, yielding the local weights w_{ilc} .

Where more than one expert is consulted, as is the case for the elicitation described in Section 3.5, each expert's judgements are first recorded as a separate matrix and only subsequently combined: for every pairwise comparison, the individual judgements are aggregated through their geometric mean, as expressed in equation (3.9), a procedure established in the AHP literature as the Aggregation of Individual Judgments method (Forman & Peniwati, 1998).

$$a_{ij}^{\text{group}} = \left(\prod_{r=1}^m a_{ij}^{(r)} \right)^{1/m} \quad (3.9)$$

where:

- m = number of experts who provided a judgement;
- $a_{ij}^{(r)}$ = judgement expressed by expert r for the pair of elements (i, j);
- a_{ij}^{group} = aggregated group judgement for the same pair, used to build the group comparison matrix from which the priority vector w is then derived with the procedure set out above.

The global weight of a decision variable, used in the scoring procedure of Section 3.3.5, is the product of the weight of the macro-category to which it belongs and its local weight

within that category, as expressed in equation (3.10). This relation follows directly from the two-level structure of the hierarchy and requires no source beyond the AHP literature already cited for that structure.

$$W_i = W_c \times w_{i|c} \quad (3.10)$$

where:

- W_i = global weight of decision variable i , referred to the hierarchy as a whole;
- W_c = weight of the macro-category c to which variable i belongs, obtained from the first-level comparison among the six categories;
- $w_{i|c}$ = local weight of variable i within category c , obtained from the second-level comparison among the variables belonging to that category alone.

3.3.3 Consistency Check

A pairwise comparison matrix built from human judgement is not guaranteed to be perfectly transitive: an expert may, for instance, judge element i more important than j , j more important than k , and yet k more important than i , a cyclical contradiction that undermines the reliability of the priority vector derived from the matrix. The Analytic Hierarchy Process addresses this through the Consistency Index, computed from the principal eigenvalue of the matrix, as expressed in equation (3.11) (Saaty, 1980, 2008).

$$CI = \frac{\lambda_{\max} - n}{n - 1} \quad (3.11)$$

where:

- n = size of the comparison matrix, that is, the number of elements compared at that level of the hierarchy.

The Consistency Index is then compared against the average Consistency Index that would be obtained from a large number of randomly generated reciprocal matrices of the same size, the Random Index, tabulated by Saaty as a function of n and reproduced in Tab. 3.2. The ratio of the two, the Consistency Ratio, is expressed in equation (3.12), and a matrix is conventionally regarded as sufficiently consistent when it falls below the threshold given in equation (3.13) (Saaty, 1980, 2008). For matrices of size $n \leq 2$, the Random Index is zero and the Consistency Ratio is undefined. Such matrices are, however, consistent by construction, since a single judgement and its reciprocal cannot contradict each other, and they are therefore exempt from the check.

$$CR = \frac{CI}{RI} \quad (3.12)$$

where:

- RI = Random Index, the average Consistency Index of a large number of randomly generated reciprocal matrices of size n , tabulated as follows.

n	1	2	3	4	5	6	7	8	9	10
RI	0.00	0.00	0.58	0.90	1.12	1.24	1.32	1.41	1.45	1.49

Tab. 3.2 - Random Index (RI) as a function of matrix size n (Saaty, 1980)

$$CR \leq 0.10 \quad (3.13)$$

Consistent with the sequencing set out in Section 3.3.2, consistency is checked at two points rather than one. Each expert's individual matrix, at both levels of the hierarchy, is checked

for consistency first. Where an individual matrix exceeds the adopted consistency threshold, the corresponding pairwise judgements are reviewed with the expert and revised where necessary, after which the consistency ratio is recalculated. Once the individual matrices meet the adopted consistency threshold, the judgements are aggregated through the geometric mean of equation (3.9) into the group matrix. The Consistency Ratio is then computed a second time on the aggregated group matrix itself, and it is this second check, performed on the matrix that generates the weights used by the DSS, that accompanies the reported weights. Should the group matrix exceed the threshold, the procedure returns to the individual judgements at the corresponding level, which are reviewed with the experts before the aggregation is repeated, since re-aggregating unchanged judgements would reproduce the same group matrix. Throughout this thesis, the Consistency Ratio is presented strictly as a check on the internal logical coherence of the judgements underlying a matrix, that is the absence of cyclical contradictions of the kind illustrated above, and not as evidence that the substantive content of those judgements is correct. The two are conceptually distinct, and satisfying the former does not by itself establish the latter (Saaty, 2008).

3.3.4 CBM Performance Matrix and Numerical Operationalisation

The decision variable×CBM matrix constructed in Section 3.2.6 records, for every decision variable i and CBM archetype j , a qualitative rating $L_{i,j}$ drawn from the domain {Low, Medium, High}, together with a benefit/cost classification δ_i indicating whether a higher rating is favourable (benefit) or unfavourable (cost) to the archetype's attractiveness. Converting this matrix into a numerical performance matrix usable by the weighted scoring procedure of Section 3.3.5 proceeds in three steps, moving deliberately from the original qualitative assessment to its directional harmonisation, and only then to a single, common numerical scale.

The first step simply encodes the original ordinal rating $L_{i,j}$ as an integer, without yet taking benefit/cost classification into account, as expressed in equation (3.14).

$$o_{i,j} \in \{1,2,3\} \quad (3.14)$$

where:

- $o_{i,j}$ = ordinal level observed for decision variable i and CBM archetype j (1 = Low, 2 = Medium, 3 = High), prior to any harmonisation.

The second step harmonises the observed level into a performance level, using the benefit/cost classification δ_i established once for each variable in Section 3.2.6. For a benefit-type variable the performance level coincides with the observed level, while for a cost-type variable it is reversed, so that a higher performance level always denotes a condition more favourable to the archetype under assessment, regardless of whether the underlying variable is of the benefit or the cost type (Hwang & Yoon, 1981). This harmonisation is expressed in equation (3.15).

$$h_{i,j} = \begin{cases} o_{i,j} & \text{if } \delta_i = \text{benefit} \\ 4 - o_{i,j} & \text{if } \delta_i = \text{cost} \end{cases} \quad (3.15)$$

where:

- $h_{i,j}$ = harmonised performance level, in which a higher value always denotes a more favourable condition for CBM archetype j with respect to decision variable i , independently of the variable's original benefit/cost classification.

The third and final step applies a single numerical scale, identically, to every harmonised level, as expressed in equation (3.16).

$$P_{i,j} = \frac{h_{i,j}}{3} \quad (3.16)$$

where:

- $P_{i,j}$ = normalised numerical performance score of CBM archetype j with respect to decision variable i , used as input to the weighted scoring formula of Section 3.3.5.

The specific numerical values 1/3, 2/3 and 1, and the assumption that the three harmonised levels are equally spaced along the resulting scale, are a modelling choice made for the purposes of this thesis: no external source attributes these particular values, and they are declared here as such rather than as an empirically derived scale. The benefit/cost distinction that motivates the harmonisation step itself, by contrast, is grounded in the established multi-attribute decision-making literature (Hwang & Yoon, 1981). This deliberately asymmetric sourcing of the two inputs to the scoring procedure, expert-elicited weights on one side and literature-derived performance ratings on the other, mirrors a distinction Saaty (1986) draws within the AHP itself between relative measurement, in which alternatives are compared pairwise as the criteria are, and absolute measurement, in which each alternative is instead rated independently against a fixed set of intensity categories. The Low/Medium/High scale developed in this section operationalises this second mode: criterion importance is elicited through pairwise comparison, while archetype performance is assessed through rating. Section 4.5.1 returns to this distinction once the DSS architecture is complete. Applying equations (3.14) to (3.16) to every cell of the decision variable $\times$ CBM matrix yields the CBM performance matrix P , with one normalised score $P_{i,j}$ for every decision variable and every archetype. The numerical mapping of equation (3.16) is applied only after the benefit/cost harmonisation of equation (3.15), so that a higher value of $P_{i,j}$ consistently represents more favourable performance regardless of whether the underlying decision variable was originally of the benefit or the cost type. The mapping does not alter the ordinal meaning of the original Low/Medium/High rating, it only operationalises that rating for the additive scoring procedure of Section 3.3.5.

3.3.5 Weighted Scoring and CBM Ranking

With the global weights W_i obtained in Section 3.3.2 and the performance matrix P obtained in Section 3.3.4, the overall score of CBM archetype j is computed as a weighted sum of its performance scores across all decision variables, following the Simple Additive Weighting logic for aggregating weights and performance scores in a multi-criteria decision problem (Kaliszewski & Podkopaev, 2016), as expressed in equation (3.17).

$$S_j = \sum_{i=1}^n W_i P_{i,j} \quad (3.17)$$

where:

- S_j = overall score of CBM archetype j ;
- n = total number of decision variables in the hierarchy;
- W_i = global weight of decision variable i , obtained in Section 3.3.2;
- $P_{i,j}$ = harmonised and normalised performance score of CBM archetype j with respect to decision variable i , obtained in Section 3.3.4.

The archetype recommended by the DSS for a given decision context is the one that maximises this overall score, as expressed in equation (3.18).

$$j_{\text{opt}} = \operatorname{argmax}_j S_j \quad (3.18)$$

where:

- j ranges over the set of the three archetypes {Product-oriented, Use-oriented, Result-oriented}.

The global weights W_i are fixed once obtained (Section 3.3.2), while the performance scores $P_{i,j}$ vary with the scenario; equation (3.17) is the mechanism through which the analysis in Section 3.4 operates: a scenario changes some of the $P_{i,j}$ inputs, and equation (3.18) is then re-applied to establish whether the recommended archetype j_{opt} changes as a result.

3.4 Scenario Analysis Approach

The weighted scoring procedure developed in Section 3.3 produces, for a given decision context, a single overall score for each of the three CBM archetypes and a single recommended archetype. That procedure, however, is only ever as informative as the performance matrix on which it operates, and some of the decision variables feeding that matrix respond causally to conditions of the external context, as established in Section 3.2.3. This section develops the mechanism through which one such exogenous condition, the state of critical raw material (CRM) markets, is used to test whether the recommendation produced in Section 3.3 holds across plausible variation in that condition or is instead sensitive to it. It proceeds in two steps: it first defines the three qualitative CRM market scenarios used for this purpose (Section 3.4.1) and then sets out the procedure through which the weighted scoring formula of Section 3.3.5 is re-applied under each of them (Section 3.4.2).

3.4.1 Definition of CRM Market Scenarios

The relevance of CRM market conditions as the exogenous condition against which the DSS recommendation is tested follows directly from the role attributed to the OEM throughout this thesis: an actor whose exposure to upstream material markets and whose incentives to adopt circularity strategies, are affected by the supply security and price conditions of the CRMs embedded in its products (Cimprich et al., 2023). At firm level, exposure to critical material supply disruptions can translate into increased production costs and broader implications for firm performance (Griffin et al., 2019). CRM market conditions are therefore treated here as an exogenous uncertainty: a feature of the external environment that lies outside the OEM's control and whose alternative states may affect the relative attractiveness of the strategic options under evaluation (Schoemaker, 1995). Framed in this way, CRM market conditions constitute a condition of the decision context that is common to all three CBM archetypes at any given point in the analysis.

Two scenario deviations are defined relative to the baseline established in Sections 3.2 and 3.3 for the purposes of this scenario analysis and are referred to throughout as the Low and High CRM Market Scenarios. The baseline itself represents a neutral, intermediate CRM market condition rather than a claim about current market conditions, and both deviations are defined relative to it. Each scenario combines two conceptually distinct dimensions of market conditions: CRM price level and price predictability. The Low scenario represents a lower-price and more predictable CRM environment, whereas the High scenario represents

a higher-price and less predictable environment. The combination of these dimensions is a deliberate scenario-design choice rather than an assumption that price level and predictability necessarily move together; the two dimensions may in principle vary independently. The scenarios therefore represent contrasting joint market states rather than measures of statistical volatility alone.

The Low and High scenarios are defined as ordinal states rather than as numerical price bands. This follows the reasoning already set out in Section 3.2.5 for the decision variables themselves. An ordinal characterisation does not require data that are rarely available in directly comparable form, while avoiding the need to specify numerical thresholds that would be difficult to justify across different CRM markets. The qualitative basis for the scenarios is provided by the critical-material literature, which documents firms' exposure to supply disruptions, price volatility and associated cost pressures across CRMs (Gaustad et al., 2018; Griffin et al., 2019). The use of a small number of contrasting scenarios to examine the implications of uncertain external conditions of this kind is consistent with the scenario-planning literature, in which scenarios are used to explore alternative plausible conditions rather than to predict a single future state (Schoemaker, 1995).

These scenarios are therefore not intended as a market forecast, nor as a claim that CRM market conditions will in fact move to one of exactly two alternative states. They are methodological constructs used to test the DSS recommendation under controlled variation in an exogenous condition that the OEM does not control. Consistent with the demonstration-oriented evaluation of the Design Science Research artefact described in Section 3.1, their purpose is to probe how the artefact behaves under contrasting illustrative market conditions rather than to generalise statistically about future CRM markets.

3.4.2 Scenario-Based Sensitivity Procedure

The scenario-based sensitivity procedure re-applies the weighted scoring formula of Section 3.3.5, equation (3.17), separately for each of the two CRM market scenario deviations defined in Section 3.4.1, introducing a scenario index s alongside the archetype index j and the decision variable index i , as expressed in equation (3.19). The baseline configuration itself is not re-computed under this equation: its score is already given by equation (3.17), since the baseline is the configuration from which equation (3.17) was originally computed in Section 3.3.5.

$$S_{j,s} = \sum_{i=1}^n W_i P_{i,j,s} \quad (3.19)$$

where:

- $S_{j,s}$ = overall score of CBM archetype j under scenario s ;
- n = total number of decision variables in the hierarchy, as in equation (3.17);
- W_i = global weight of decision variable i , obtained once in Section 3.3.2;
- $P_{i,j,s}$ = harmonised and normalised performance score of CBM archetype j with respect to decision variable i , under scenario s .
- s = index of the scenario deviation, $s \in \{\text{Low, High}\}$, as defined in Section 3.4.1.

Equation (3.19) differs from equation (3.17) only in the third index carried by the performance term. The global weights W_i remain fixed (Section 3.3.2), isolating the effect of changing market conditions on the performance of the alternatives. What is allowed to

change from one scenario deviation to another is the performance score $P_{i,j,s}$ itself, and only for the subset of decision variables for which a causal link to CRM market conditions can be established. This scenario-sensitivity condition is expressed formally through a binary indicator c_i , as expressed in equation (3.20).

$$P_{i,j,s} = \begin{cases} P_{i,j}^{(s)} & \text{if } c_i = 1 \\ P_{i,j} & \text{if } c_i = 0 \end{cases} \quad (3.20)$$

where:

- c_i = binary indicator equal to 1 if decision variable i is causally linked to CRM market conditions, and 0 otherwise;
- $P_{i,j}^{(s)}$ = performance score of CBM archetype j with respect to scenario-sensitive decision variable i , reassessed under scenario deviation s ;
- $P_{i,j}$ = baseline performance score of decision variable i established in Section 3.3.4, retained unchanged when $c_i = 0$.

Decision variables for which no causal link can be established correspond to $c_i = 0$ and retain their baseline performance score across both scenario deviations.

For each scenario deviation, the archetype recommended by the DSS is again the one that maximises the corresponding overall score, as expressed in equation (3.21).

$$j_{\text{opt},s} = \underset{j}{\operatorname{argmax}} S_{j,s} \quad (3.21)$$

where:

- $j_{\text{opt},s}$ = archetype recommended by the DSS under scenario s , ranging over the set of the three archetypes {Product-oriented, Use-oriented, Result-oriented}, as in equation (3.18).

Applying equation (3.21) to the Low and High scenario deviations, alongside the baseline ranking already obtained from equation (3.18) in Section 3.3.5, yields three recommendations, which may point to up to three different archetypes: one for the baseline configuration and one for each of the two scenario deviations. If all three coincide, the recommendation produced in Section 3.3 is robust to variation in this exogenous condition. If they diverge, the recommendation is sensitive to it and the specific pattern of divergence, which archetype is favoured under which deviation from the baseline, becomes itself a substantive finding of the analysis rather than an inconclusive result. Both outcomes are treated as legitimate and informative, since no pattern of convergence or divergence is assumed in advance.

Moving from one scenario to another only requires updating the scenario-dependent performance scores $P_{i,j,s}$ and re-applying equations (3.19) and (3.21). The scenario analysis is, in this sense, a computationally light extension of the weighted scoring procedure already developed in Section 3.3.

3.5 Interview Design and Data Collection

This section describes the expert interviews conducted for this research. The interviews served two purposes. The first was to obtain a brief expert review of the framework developed in Section 3.2, focusing on the clarity, relevance and perceived completeness of its categories and decision variables. The second was to collect the pairwise judgements required to calculate the AHP weights described in Section 3.3. The expert review was not intended to generate the decision variables, which had been defined before the interviews,

but to provide an additional qualitative check of the framework. Section 3.5.1 presents the interview protocol, while Section 3.5.2 describes the expert selection process. Consistent with the relevance cycle described in Section 3.1 (Hevner, 2007), the interview connects the DSS with individuals who have relevant experience in offshore wind and in the decision problem addressed in this thesis. The interview protocol is provided in Appendix A.

3.5.1 Expert Review and Structured AHP Elicitation

The interview followed a guide prepared before the sessions, covering both an open, exploratory part and a structured comparison exercise. Following recommended practice for semi-structured interview design (Kallio et al., 2016).

Four interviews were conducted, three online via Microsoft Teams and one in person, each lasting approximately one hour. In both modes, the interviewee was provided with the same guiding document, shared on screen during the online sessions and provided in printed form during the in-person session. The interviewer then took each interviewee through the document step by step. The material shown, question order, and overall session structure were kept constant across delivery modes, reducing procedural variation between interviews. All interviews were audio-recorded and initially transcribed using Otter.ai. The transcripts were subsequently checked against the original recordings and corrected where necessary.

Before each interview, the interviewee was informed that participation was voluntary, that the session would be recorded, and that any information provided would be treated confidentially and used for academic purposes only. Recording began only after the interviewee confirmed that they were comfortable with this. Interviewees were also given the choice of being identified by name and professional role or having their contribution anonymised. The profiles reported in Tab. 3.3 reflect the preference expressed by each interviewee. This approach is consistent with the literature emphasising the explicit consideration of confidentiality and anonymity in research design (Wiles et al., 2008).

The first part asked about the interviewee's background and relevant experience and included a brief review of the framework developed in Section 3.2. Participants were asked whether the six-category hierarchy and decision variables were clear, relevant to the OEM's decision problem, and sufficiently comprehensive, and whether any category appeared to overlap with another or any relevant factor was missing. No participant suggested an additional decision variable for inclusion in the framework. This review was used as a qualitative check of the framework's clarity, relevance and perceived completeness and was kept distinct from the AHP judgements collected in the second part.

The OEM is the reference decision maker for the DSS. The comparison exercise therefore asked each expert to assess the relative importance of the criteria for this reference decision problem, rather than according to their own institutional interests. This kept the object of judgement consistent across experts with different backgrounds. It reflects a broader point in the MCDA literature: the decision maker whose problem is being modelled is not necessarily the person who provides the preference information used to model it (Cinelli et al., 2020). Eliciting this information from selected experts, rather than directly from the decision maker, is one established approach to gathering the preference information required by MCDA (Marttunen et al., 2015). As none of the interviewees was employed by an OEM, their judgements were treated as informed proxy assessments of the reference decision

problem rather than as direct measurements of OEM preferences. This applies also to interviewees with first-hand experience of commercial interactions with OEMs.

The second part reproduced the AHP hierarchy presented in Section 3.3. The six macro-categories were compared pairwise, resulting in 15 comparisons for $n = 6$, followed by pairwise comparisons of the decision variables within each category, using Saaty's 1-9 scale (Saaty, 1980, 2008). A short example unrelated to the thesis topic preceded the first comparison to allow the interviewee to practise using the scale. The comparison order was fixed in advance and kept identical across interviews. The interviewer recorded the selected option and intensity for each pair during the session. The full reciprocal matrices, priority vectors, and consistency ratios were derived afterwards.

Because the consistency ratio (CR) can only be computed once a complete comparison matrix is available, it was calculated after each interview rather than during the session. The resulting priority vectors and CR were then shared with the interviewee, together with a summary of the information gathered during the session. This gave each expert an opportunity to confirm that their recorded judgements had been captured accurately. Sharing findings with participants in this way is a recognised means of strengthening the credibility of interview data (Birt et al., 2016). Where the CR of a comparison matrix exceeded the threshold specified in Section 3.3, the relevant pairwise judgements were reviewed with the interviewee. They were asked to reconsider and, where appropriate, revise the comparisons that appeared to contribute to the inconsistency. This approach retained the expert as the source of the revised judgement rather than imposing a mathematical adjustment or discarding the original judgement, consistent with guidance on handling inconsistency in AHP applications (Ishizaka & Labib, 2011). Both the original and revised judgements were retained in the research record.

The interview closed with a second completeness check. After completing all comparisons, the interviewee was asked again whether any relevant decision variable was still missing. No participant identified an additional variable for inclusion and no final ranking or aggregate result was shown at this stage.

3.5.2 Expert Selection and Sampling

Interviewees were selected purposively based on demonstrated professional or research experience relevant to the OEM's decision problem, consistent with the general logic of expert sampling: participants are selected because their experience makes them likely to provide informed and relevant judgements (Patton, 2002).

Selection focused on demonstrated professional or research experience relevant to the OEM decision problem in the offshore wind sector. Relevant experience could concern different dimensions of the decision problem, including technical, commercial and economic, organisational, regulatory, market or supply-chain aspects. Familiarity with servitised or "as-a-service" CBMs was considered an additional desirable characteristic rather than a strict requirement. Candidates were identified through academic and industry networks and approached individually. Tab. 3.3 summarises the profile of each interviewee.

ID	Interviewee	Profile	Category	Relevant experience
E1	Matthew Chapman	Marketing Manager, Natural Power; previously Marketing Manager, Renewable Parts Limited	Informed proxy (consultant/industry advisor)	11 years in renewable energy, of which approximately 5 specific to offshore wind; cross-sector work on circular economy policy, skills and end-of-life strategy
E2	Colin Andrews	Retired academic (operations management); PhD on the emergence of offshore wind as an industry; involved in establishing the Scottish Institute for Remanufacturing	Informed proxy (academic)	≈15 years studying the offshore wind sector; longitudinal PhD research (2011-2020) on industry emergence; early remanufacturing work with Renewable Parts Limited
E3	Remi Zante	Chief Engineer, Renewable Parts Limited; also affiliated with the National Manufacturing Institute Scotland (NMIS)	Informed proxy (Supply-chain practitioner)	≈3 years specific to offshore wind; background in manufacturing across several sectors; direct involvement in component refurbishment/ remanufacturing, currently undergoing the audited-supplier assessment process with OEMs

ID	Interviewee	Profile	Category	Relevant experience
E4	Paul McManus	Procurement Manager, ScottishPower Renewables	Informed proxy (Buyer-side practitioner, direct OEM counterpart)	≈3 years specific to wind-sector procurement, within a broader career in procurement and supply chain management (chartered professional); direct engagement with OEMs on contracts, pricing and asset purchasing/maintenance decisions; active sourcing of remanufactured components, critical raw materials and second-hand parts for an ageing asset base

Tab. 3.3 - Profile of the expert interviewees.

Colin Andrews contributes a research-based perspective grounded in long-term study of the offshore wind sector and experience in remanufacturing. As he is not an OEM employee, his judgements are treated as informed proxy assessments of the reference decision problem rather than as OEM-insider judgements.

Paul McManus contributes a buyer-side perspective based on direct commercial interaction with turbine OEMs, including negotiations concerning contracts, pricing, purchasing and maintenance decisions. This provides first-hand exposure to commercial interactions with OEMs, while, as he is not employed by an OEM, his judgements are treated as informed proxy assessments rather than direct measurements of OEM preferences.

The panel combines different types of relevant expertise, providing some variation in the sources of contextual knowledge while keeping all judgements focused on the same OEM decision problem (Section 3.5.1). It is an exploratory expert-elicitation panel rather than a statistically representative sample. Its purpose is to obtain informed, traceable judgements and to populate and demonstrate the DSS in the illustrative case presented in Chapter 5, not to estimate the preferences of a wider population of OEM decision-makers. Individual judgements were combined using geometric-mean aggregation, following the established treatment of group judgements in AHP (Forman & Peniwati, 1998). Forman and Peniwati (1998) distinguish this procedure, termed Aggregation of Individual Judgments (AIJ), from the alternative of first deriving each expert's own priority vector and only then aggregating the resulting vectors (Aggregation of Individual Priorities, AIP) and recommend the former when the group is expected to act as a single, synergistic decision-making body rather than as an assembly of separate stakeholders whose individual priorities should be preserved and combined. This condition fits the panel assembled for this thesis: the four experts were not

asked to represent distinct institutional interests in competition with one another, but to converge, individually, on an informed assessment of the same reference decision problem (Section 3.5.1). Aggregating judgements before deriving priorities, rather than aggregating already-derived priorities, was therefore the procedure consistent with the purpose for which the panel was constituted. The limited panel size constrains the generalisability of the resulting weights beyond the four experts interviewed.

4. Framework Development: From CBM Value-Chain Mapping to the Decision Support System

4.1 Value Chain / Actor Mapping of the Three CBM Archetypes

This section maps three CBM archetypes based on product-service systems (PSS): product-oriented, use-oriented and result-oriented offers (Tukker, 2004; Baines et al., 2007). This mapping applies the circularity lens introduced earlier in this thesis, which distinguishes between an archetype's structural control over whether and how an asset returns for reuse and the technical feasibility of recovering the material it contains once it does return. This section, and the framework built on it, addresses only the first of these two dimensions. Technical recovery feasibility is therefore treated as a technology-dependent condition rather than as a property of the CBM archetype itself. This distinction is important when interpreting the framework, because structural control over the return of an asset does not in itself imply that the material contained in it can be technically recovered. Technical recovery feasibility may therefore remain a critical condition for material recovery, but it is assessed separately from the CBM archetype structure. OEM designates the manufacturer, treated throughout as the single actor that both originates and delivers the offer.

Two sets of conditions bound this mapping. The first applies across the whole chapter: the OEM operates in a business-to-business, capital-intensive setting in which the asset under consideration is durable and contains critical raw materials, and it is the OEM, rather than the customer or a third party, that chooses among the three archetypes evaluated here. These conditions carry forward the industrial context already established earlier and are therefore not re-argued in this chapter. The second concerns the three archetypes individually. Comparing the three CBM archetypes on a consistent basis requires fixing a single reference configuration for each one, rather than attempting to represent every possible implementation at once.

Under the product-oriented configuration, the OEM sells the asset outright and ownership passes to the customer at the point of sale. Lifecycle services such as maintenance or take-back may still be offered and separately contracted, but they remain optional additions to the transaction rather than structural features of it. For rating purposes in Sections 4.3 and 4.4, the product-oriented OEM is assumed to offer a take-back mechanism for its products, in addition to any statutory obligation. Whether, in what quantity and in what condition products actually return remains dependent on the customer's decision.

Under the use-oriented configuration, the OEM retains ownership of the asset and grants the customer access to it rather than title, capturing value through continuous payments over the life of the arrangement rather than a single price at delivery. The OEM also retains responsibility for maintenance and technical control of the asset throughout. The customer's role is limited to deciding how, how much and for how long to operate the asset within the terms of the agreement.

Under the result-oriented configuration, the OEM retains ownership of the underlying asset and assumes contractual responsibility for delivering an agreed functional or performance result. The OEM also retains decision authority over the technical means, maintenance and lifecycle decisions concerning that asset. The customer may, however, retain operational

responsibility for the wider system within which the result is delivered, without itself operating the asset through which the OEM delivers that result.

Each archetype is first described individually; Section 4.2 then compares them. Each value chain is unpacked across three life-cycle phases: beginning-of-life (BOL) covering design, sourcing, production and contracting; middle-of-life (MOL) covering distribution, use, monitoring and maintenance; and end-of-life (EOL) covering collection, recovery, remanufacturing, recycling and disposal (Gaustad et al., 2018; Govindan et al., 2015). For each phase, this section reports the actors involved, the activity each performs, and whether the OEM controls that activity directly or depends on another actor.

4.1.1 Product-Oriented CBM

The OEM's value proposition is the sale of a physical product, complemented where relevant by product-related services such as maintenance contracts (Tukker, 2004). The OEM captures value through the sale price and, where contracted, through payment for these services (Reim et al., 2015). Ownership of the product, including its constituent materials, passes to the customer at the point of sale, consistent with the product-oriented reference configuration set out above. Any residual influence the OEM retains afterwards therefore depends on what was contracted, rather than on the ownership structure of the model itself (Reim et al., 2015).

Beginning of life. The actors are the OEM's design and engineering function, upstream suppliers and, where a distribution layer exists, dealers or intermediaries (Reim et al., 2015). Because the transaction is settled and paid upfront at the point of sale, this archetype does not require the OEM to structurally finance continued customer access to the asset over the contract life, unlike the two ownership-retaining archetypes (Tukker, 2004). Where a service or take-back element is present, it is negotiated as an optional addition to the transaction rather than as a structural feature of the model (Tukker, 2004; Lewandowski, 2016).

Middle of life. The customer is the sole actor deciding how, how much and for how long to operate the product. Maintenance may be delegated either to a third party or to the OEM (Tukker, 2004; Reim et al., 2015). The OEM's information about the product in use is limited to whatever has been specifically contracted (Yang et al., 2018).

End of life. Where no take-back mechanism exists, the customer returns the product, if at all, to an independent scrap dealer, reseller or waste-management operator. Where a take-back mechanism exists, the product returns to the OEM, potentially through a reverse-logistics provider (Gaustad et al., 2018; Govindan et al., 2015). Under an individual scheme, the OEM remains in control of the units it collects. Under a collective scheme, exemplified by take-back obligations under Waste Electrical and Electronic Equipment legislation for categories such as professional electrical and electronic equipment, this control is shared with regulatory bodies or scheme administrators (Gaustad et al., 2018; Pollard et al., 2023). Even under an individual scheme, the quantity and quality of returns remain a source of uncertainty (Govindan et al., 2015). Moreover, a returned product may still fail to be recovered if it was not designed for recovery (Pollard et al., 2023).

4.1.2 Use-Oriented CBM

Ownership of the product stays with the OEM, consistent with the use-oriented configuration set out above; the customer is granted access rather than title (Tukker, 2004). The OEM captures value through continuous payments over the life of the arrangement rather than through a single price at delivery (Reim et al., 2015).

Beginning of life. Actors extend beyond the OEM's design function to include financial institutions that fund a revenue stream collected over the life of the arrangement rather than at delivery (Reim et al., 2015). Because the same unit is expected to serve multiple contract periods or users, design requirements extend beyond ease of maintenance to include durability, upgradability and re-manufacturability (Reim et al., 2015). This is consistent with the use-oriented CBM archetype, in which the party granting access also assumes responsibility for service and maintenance (Bocken et al., 2016). The offer further requires infrastructure for sensing, connectivity and data platforms to locate the asset and monitor its condition and use (Hidalgo-Crespo et al., 2025). Because ownership stays with the OEM throughout, the resulting incentive for durable design reaches the same actor that makes the corresponding design decisions (Tukker, 2004).

Middle of life. Responsibility for maintenance and control during use rests with the OEM, either directly or through third-party service providers it engages (Tukker, 2004). This requires closer monitoring of the customer and the asset than a product-oriented offer, generating data that can inform product adjustments and subsequent contract terms (Reim et al., 2015; Hidalgo-Crespo et al., 2025). Because customers do not bear the full consequences of asset deterioration when they do not own the asset, this arrangement can encourage less careful use, a risk that the OEM must manage contractually (Tukker, 2015; Reim et al., 2015).

End of life. Because legal ownership never transfers, the OEM does not need to negotiate a return right after the fact, as a product-oriented OEM does. Recovery is instead structurally incorporated into the lease arrangement through the physical hand-back of the asset once the contract term ends, with the return handled through the OEM's own maintenance and reverse-logistics network (Yang et al., 2018). Because the OEM is incentivised to extend the asset's life and recapture value from it, this network can support multiple cycles of reuse, remanufacturing and recycling over the asset's operating life (Yang et al., 2018). This potential is realised only where the party managing the asset at end of life is also positioned to act on the original design (Tukker, 2004).

4.1.3 Result-Oriented CBM

The OEM and customer agree on a result rather than on a predetermined product, consistent with the result-oriented configuration set out above (Tukker, 2004). Ownership of the underlying asset remains with the OEM throughout, while the customer's claim is limited to the paid-for result (Baines et al., 2007; Reim et al., 2015).

Beginning of life. On the OEM's side, actors include specialised subcontractors for tasks beyond its own capabilities, such as financial institutions and recycling or transport companies (Reim et al., 2015). On the customer's side, the archetype typically involves direct interaction with a small number of business-to-business counterparts, with whom trust

must be established before the offer can be extended (Reim et al., 2015). Because a promised result can only be verified and priced if its delivery can be observed, the network also extends to providers of sensing and data-platform infrastructure through which that delivery is confirmed (Hidalgo-Crespo et al., 2025). Materials and consumables are cost items rather than sources of revenue, giving the OEM a direct financial incentive to reduce material use and reuse components where design permits. Realising this incentive, however, requires redesigning delivery around the underlying product rather than only relabelling an existing offer as a service (Tukker, 2015).

Middle of life. The result must be delivered continuously over the contract, sustained by ongoing contact between the OEM's operating personnel and the customer rather than by scheduled visits (Reim et al., 2015). This generates comprehensive information on asset performance, which feeds back into beginning-of-life design and sourcing decisions (Reim et al., 2015; Hidalgo-Crespo et al., 2025). Because the underlying assets, monitoring infrastructure and components typically originate from different suppliers, the OEM coordinates a wide network of external partners while retaining decision authority over the technical means and lifecycle management of the asset throughout (Hidalgo-Crespo et al., 2025). In exchange, it assumes the operational and financial risk of delivering the result at a fixed, pre-agreed price. This risk can only be assumed by an OEM with the capacity to absorb it (Reim et al., 2015; Tukker, 2015).

End of life. Because the OEM never relinquishes the asset, there is no ownership-transfer or hand-back event to negotiate, unlike under the use-oriented configuration. Recovery is instead embedded in the OEM's continuing management of the result-delivery system (Yang et al., 2018). Because operational savings accrue directly to the OEM, the incentive to minimise resource use is correspondingly strong, and the design is not tied to any pre-existing product concept (Reim et al., 2015). Realising this potential depends on the OEM's capacity to absorb a risk premium and coordinate an external ecosystem of monitoring, data and component partners (Tukker, 2004).

4.2 Cross-Model Comparison: From Value Chains to Candidate Decision Factors

4.2.1 Structural Divergence across the Three Value Chains

Beginning of life. The most consequential divergence concerns who holds property rights to the asset once it leaves production. The product-oriented archetype settles the transaction at sale and releases the OEM from further capital exposure, whereas both ownership-retaining archetypes require the OEM to fund and continue funding an asset it does not sell (Tukker, 2004; Reim et al., 2015). This exposure is not uniform even between the two: it is higher under the result-oriented configuration, whose OEM also absorbs the liability of delivering the promised result, than under the use-oriented one (Tukker, 2004; Tukker, 2015). This capital question is distinct from the ordinary cost of delivering the offer in terms of labour, process and field service. This cost also rises across the three archetypes, but for a different reason: it reflects operational responsibility once the asset is in use rather than the balance-sheet commitment made at sale (Reim et al., 2015; Tukker, 2015).

A related divergence, fixed at the same point of origination, concerns the agreement itself. A product-oriented transaction is highly formalised and low in complexity and leaves little risk

with the OEM once concluded. A use-oriented contract carries proportionately more of both as ownership is retained, while a result-oriented agreement is the most complex and least standardised of the three, since its performance criteria and risk allocation must be negotiated afresh for each engagement (Reim et al., 2015).

Ownership retention also creates a direct financial stake in durable, recovery-oriented design that a point-of-sale transaction does not (Tukker, 2015; Bocken et al., 2016). This is distinct from the freedom to design: a product-oriented offer ties the OEM to one saleable product, whereas use- and, more markedly, result-oriented offers widen the range of technical solutions available, since the customer contracts for use or for an outcome rather than for a specific product (Tukker, 2004). Neither gradient translates into practice on its own. Exploiting it requires a shift from a culture built around unit sales to one built around a sustained service relationship, a transformation the servitisation literature treats as substantial rather than incremental and that structured readiness assessments increasingly evaluate against a company's existing capabilities (Baines et al., 2009; Frederiksen et al., 2021). A final beginning-of-life divergence concerns the intermediary network the OEM must build to reach the market: conventional dealers for a product-oriented sale, financing partners for both ownership-retaining archetypes, and, for the result-oriented offer specifically, more direct coordination with customers and specialised external partners (Tukker, 2004; Reim et al., 2015).

Middle of life. The clearest divergence is the extent of visibility the OEM retains into how the asset is used. This moves from a limited, contract-bound data flow under product-oriented offers to the comprehensive and continuous monitoring obtained by a result-oriented OEM through its continuing decision authority over the asset's technical means and lifecycle management, with use-oriented offers falling in between (Reim et al., 2015; Yang et al., 2018). Building the infrastructure that produces this visibility is a separate and cumulative demand, growing from product- to result-oriented (Baines et al., 2007; Hidalgo-Crespo et al., 2025). Who then controls and can act on the resulting data is a governance question layered on top of it. Maintenance responsibility follows the same gradient, shifting from the customer, or a third party it engages, to the OEM as the archetype moves from product- to result-oriented (Tukker, 2004; Reim et al., 2015). A related but distinct divergence concerns the intensity of customer contact itself: episodic under a product-oriented offer, closer under a use-oriented one, and sustained through continuous contact under a result-oriented offer (Reim et al., 2015).

A further divergence concerns customer behaviour: a customer who does not own the asset has less incentive to care for it (Tukker, 2015). Under a result-oriented arrangement, the OEM's retained decision authority over the asset's technical means and maintenance limits the customer's opportunity to under-care for the core asset itself. Exposure to that particular risk therefore falls again below the use-oriented case, even though the customer's residual responsibility for the wider system within which the result is delivered still carries a documented, if smaller, non-ownership care risk of its own (Yang et al., 2018; Tukker, 2015). A final divergence concerns payment timing. A product-oriented sale realises its revenue once, largely independent of subsequent use, while use- and result-oriented offers collect it over the contract's life, tied to continued access or delivered outcome (Reim et al., 2015). This makes the OEM's economics considerably more dependent on that demand materialising as forecast.

End of life. The nature of the return event itself changes across the three archetypes, from a conditional right that the OEM must negotiate and hope the customer exercises, to a structural feature of a lease term, and finally to a condition that does not arise where the OEM never relinquishes the asset (Yang et al., 2018; Gaustad et al., 2018). Whether returned material can be monetised is a distinct question, depending on downstream capabilities such as remanufacturing capacity or access to a market for recovered material, and is largely independent of the archetype's ownership structure (Cimprich et al., 2023).

A further end-of-life divergence concerns the coordination burden of securing the return, which does not decline monotonically with ownership retention. A product-oriented OEM depends on reverse-logistics and recycling infrastructure it typically does not own, and the quantity and quality of what returns remain uncertain (Govindan et al., 2015). A use-oriented OEM works through an established, comparatively standardised network of third-party partners built around its own asset base (Reim et al., 2015). A result-oriented OEM internalises control of the core asset but depends on a wider constellation of specialised subcontractors for monitoring, data and components (Hidalgo-Crespo et al., 2025). The burden is therefore high at both ends of the ownership spectrum, for structurally opposite reasons.

An OEM weighing the three archetypes faces a consistent set of considerations recurring across each life-cycle phase: capital and operating costs outstanding on its side; the formalisation, complexity and risk of the agreement; the incentive, freedom and organisational capability to design well; the visibility into and interaction around the asset's use; how the customer is likely to behave once it does not own the asset; how certain and how monetisable the asset's eventual return is; and how complex the coordination network becomes. None of these considerations, however, captures a condition applying with equal force to all three archetypes, since such a condition could not emerge from a comparison built in this way.

4.2.2 Context and Scenario Conditions Beyond the Comparison

A comparison built on differences between CBM archetypes is, by design, blind to conditions that apply equally to all three. A general reluctance to consume without owning, the volatility of a critical-material market, the pace of technological or regulatory change and the strictness of the regulatory regime do not vary by archetype. Ignoring them would not sharpen the analysis but leave it incomplete, since they shape how much the structural differences identified in Section 4.2.1 matter in practice. This section draws on the same literature used above, together with a systematic review of the drivers and barriers to CBM adoption (Assmann et al., 2023) and a comparative case study of CBM innovation (Geissdoerfer et al., 2023).

Customers' underlying desire for ownership and their resistance to unfamiliar contractual arrangements are documented as persistent barriers to offers that depart from outright sale. These barriers apply to both ownership-retaining archetypes, although not necessarily to the same degree (Assmann et al., 2023). They therefore set a baseline of market acceptance that both must overcome before the structural advantages identified in Section 4.2.1 can be realised. Related but distinct is the predictability of customer demand, which is a property of the specific market rather than of the archetype (Assmann et al., 2023). A lack of customer demand is among the most frequently reported barriers to circular strategies generally

(Geissdoerfer et al., 2023). This exogenous uncertainty must be kept separate from the archetype-specific dependency on forecast demand established in Section 4.2.1. A market can be highly predictable and still leave an OEM exposed, under an ownership-retaining archetype, to a forecast that turns out to be wrong.

More consequential for this thesis is the state of the CRM market, which is subject to recurrent and severe price volatility driven by supply concentration, geopolitical exposure and clean-technology demand (Gaustad et al., 2018; Cimprich et al., 2023). This condition affects every OEM regardless of archetype. What varies, following from the return-certainty divergence of Section 4.2.1, is the extent to which recovered material can offset it. An archetype under which material rarely returns leaves the OEM dependent on virgin procurement at whatever price the market sets. An archetype under which the asset is reliably recovered can meet part of its future material needs internally, at a cost that is more insulated from that volatility. This is the mechanism that the scenario analysis developed later in this thesis is designed to examine, moving this condition through Low, Baseline and High states while the three reference configurations set out in Section 4.1 remain unchanged; it is the severity of the market condition that varies, not the archetypes themselves.

A different exogenous condition concerns the pace of technological and regulatory change relative to the length of commitment imposed by an archetype. A product-oriented sale resets this exposure at each transaction, since a new sale is priced and designed for the conditions prevailing at the time. A use- or result-oriented contract instead commits the OEM for the length of a single agreement, since the investment it requires is not easily reversed within that period (Tukker, 2015). This commitment is not, however, uniformly a commitment to a fixed technology. A use-oriented contract ties the OEM to the specific asset installed when the lease was signed, whereas a result-oriented contract commits it only to the promised outcome, leaving more room to adapt the underlying technical means during the agreement. It is the combination of long-term commitment and rapid external change, most acute where that commitment is also tied to one fixed technology, that matters most.

Regulation surrounding recovery obligations, material use and their financial support applies to an OEM regardless of archetype, but its consequences are not uniform. Take-back and extended-producer-responsibility regulation, in its individual or collective form, sets a recovery floor across all three archetypes (Gaustad et al., 2018). The product-oriented reference configuration is assumed to meet this floor through the take-back mechanism described in Section 4.1. How easy that floor is to meet depends on the archetype-specific reverse-logistics arrangements identified in Section 4.2.1. A wider set of levers, including government targets, supportive laws, taxation, public funding and industry standards, is documented as bearing on adoption more broadly and tends, by design, to favour the two ownership-retaining archetypes (Assmann et al., 2023). Regulatory pressure should not, however, be treated as a uniformly powerful driver. An independent case study found regulatory push among the least relevant drivers reported by companies actually adopting circular models, even though its absence remained one of the most frequently cited barriers (Geissdoerfer et al., 2023). An absent regulatory framework is thus a more reliable obstacle than a present one is a driver. A combination of instruments, sometimes described as a carrot-and-stick approach, is more effective than either financial incentives or binding mandates alone, since financial support alone can prove fragile once withdrawn, while mandates such

as minimum recycled-content requirements are harder to implement without sustained policy commitment (Cimprich et al., 2023).

The EU Critical Raw Materials Act illustrates how concrete this horizon already is. It sets non-binding 2030 benchmarks across seventeen strategic materials: 10% domestic extraction, 40% domestic processing and 25% recycling, together with a diversification cap of 65% of supply from any single non-member country (Regulation (EU) 2024/1252). It also sets a recycled-content trajectory relevant to the recovery divergences identified in Section 4.2.1. Manufacturers of products containing permanent magnets above a defined threshold face a public disclosure obligation from 2027, ahead of delegated acts due by 2031 that are expected to make minimum recycled-content requirements binding for certain rare-earth elements (Ragonnaud, 2024). A long-term use- or result-oriented contract signed today therefore risks encountering tighter obligations before it expires, an exposure that a product-oriented sale, renegotiated at each transaction, largely avoids. Demonstrating compliance with such obligations is, in turn, more straightforward where the OEM already tracks the asset continuously, connecting this regulatory condition to the data-visibility divergence identified in Section 4.2.1 (Pollard et al., 2023).

Everything an OEM needs to weigh when choosing among the three archetypes therefore has one of two origins: a genuine divergence between value chains, established in Sections 4.1 and 4.2.1, or a condition of the environment that applies regardless of archetype but still shapes its relative attractiveness, established in this section. Section 4.3 consolidates both into the final set of decision variables.

4.3 The Final Set of Decision Variables

Turning the pool of candidate considerations identified above into a criteria structure requires two corrections established as part of the method in Chapter 3 (Section 3.2.5): a boundary correction, which moves a candidate to the macro-category its content actually belongs to, and a construct-level consolidation, which folds two candidates that address the same underlying question into a single variable. Applying both corrections to the pool generated by Sections 4.1 and 4.2 yields eighteen decision variables: three in Market, four in Economic/Business Case, three in Technical/Product, two in Organisational/Capabilities, three in Value Chain/Ecosystem and three in Regulatory/Institutional. Each variable below records its origin, either the structural comparison in Section 4.2.1 or a context condition from Section 4.2.2, and is defined independently of any one archetype. Performance against the three archetypes is rated separately in Section 4.4.

4.3.1 Market

This category groups factors that depend on how the OEM's customers, individually and in aggregate, are disposed to respond to a proposed archetype. This disposition belongs to the customer base rather than to the OEM's own economics or technology.

Customer Acceptance of the CBM's Ownership/Access Arrangement records the customer base's general willingness to accept the ownership or access arrangement proposed by a given archetype (Assmann et al., 2023). It excludes both how a customer behaves once the arrangement is accepted, captured by the variable below, and the OEM's dependency on

demand materialising as forecast, captured by Demand/Usage Forecasting Dependency below.

Customer Asset-Care and Moral-Hazard Exposure records the extent to which the archetype exposes the OEM to reduced customer incentives to care for an asset because the customer does not own it (Section 4.2.1; Tukker, 2015; Reim et al., 2015). It sits under Market rather than Organisational/Capabilities because it measures customer-side behaviour that the OEM cannot itself change.

Demand/Usage Forecasting Dependency records the extent to which an archetype's economics depend on the volume or intensity of demand and usage materialising as forecast. It is distinct from how predictable that demand is in the market itself (Section 4.2.2; Assmann et al., 2023) and from the OEM's capacity to deliver against it, which is captured separately under Service and Maintenance Capability. This dependency is negligible for a product-oriented sale, since revenue is realised at the point of sale, and increases for the two ownership-retaining archetypes, whose revenue accrues over the contract's life rather than through a single transaction (Reim et al., 2015).

Decision variable	Description	Origin	Citations
Customer Acceptance of the CBM's Ownership/Access Arrangement	How willing, in general, the OEM's customer base is to accept the ownership or access arrangement a given archetype proposes.	Context-derived (4.2.2)	Tukker, 2015; Assmann et al., 2023
Customer Asset-Care and Moral-Hazard Exposure	The extent to which the archetype exposes the OEM to reduced customer incentives to care for an asset because the customer does not own it.	Structural comparison (4.2.1)	Tukker, 2015; Reim et al., 2015
Demand/Usage Forecasting Dependency	How far the archetype's own economic sustainability depends on demand and usage materialising as forecast.	Structural comparison (4.2.1)	Reim et al., 2015; Assmann et al., 2023

Tab. 4.1 - Final decision variables in the Market macro-category.

4.3.2 Economic/Business Case

This category groups the archetype's financial architecture along four dimensions that vary independently and must not be collapsed into a single notion of cost: investment required, ongoing operating cost, value creation and capture, and exposure to CRM price volatility. Tukker's (2004) risk-premium facet of a PSS's economics is not carried forward as a separate fifth dimension here but distributed across Provider Investment/Capital Requirement and Value Creation and Capture Potential below.

Provider Investment/Capital Requirement records how much capital and financial capacity the OEM, or a financing partner, must structurally commit and keep committed as

a consequence of the chosen archetype, net of ordinary production investment (Section 4.2.1; Tukker, 2004; Reim et al., 2015; Tukker, 2015; Frederiksen et al., 2021). This covers both capital tied up in the asset itself and the financial reserves an archetype requires the OEM to hold against the risk it assumes, such as the liability of delivering a contracted result. It is held fixed across all market scenarios tested in Section 4.4.3, because the variable records the structure of the commitment (asset ownership and the liability assumed), not its monetary amount, which would vary with material prices, while material-price exposure is captured separately below.

Operational Service-Delivery Cost Burden covers the ordinary labour, process and field-service costs of delivering the offer (Reim et al., 2015; Tukker, 2015). It excludes the capital cost of the asset and the purchase price of virgin critical raw materials, captured respectively by Provider Investment/Capital Requirement above and CRM Price Volatility Exposure below, but includes labour and service costs that depend on the choice between replacement and repair or reconditioning.

Value Creation and Capture Potential records the OEM's ability to create and capture economic value from the CBM, including transaction revenues, recurring access or result payments, service-related value and the monetisation of recovered material. It concerns the economic value that can be appropriated from the offer, not the physical return of the asset or the OEM's residual dependence on virgin CRM purchases, which are captured separately by CRM Recovery and Asset Return Effectiveness and CRM Price Volatility Exposure, respectively. A use- or result-oriented OEM is better placed here because ownership retention gives it the property-rights position from which value generated over the relationship's life can be routed back into its own chain (Cimprich et al., 2023). Why use and result are nonetheless rated equally is explained in Section 4.4.1.

CRM Price Volatility Exposure records the archetype's economic exposure to volatile virgin CRM prices, considering the extent to which future material needs can be covered by material returned under the archetype (Section 4.2.2; Gaustad et al., 2018; Cimprich et al., 2023).

Decision variable	Description	Origin	Citations
Provider Investment/Capital Requirement	Capital and financial capacity structurally tied up on the OEM's side, including reserves held against the risk the archetype assumes, net of ordinary production investment.	Structural comparison (4.2.1)	Tukker, 2004; Reim et al., 2015; Tukker, 2015; Frederiksen et al., 2021
Operational Service-Delivery Cost Burden	The ordinary labour, process and field-service cost of delivering the offer, excluding the capital cost of the asset and the purchase price of raw materials, but including labour and service costs that depend on the choice between replacement and repair or reconditioning.	Structural comparison (4.2.1)	Reim et al., 2015; Tukker, 2015
Value Creation and Capture Potential	The OEM's ability to create and capture economic value from the CBM, including recurring payments and the monetisation of recovered material, distinct from the physical return of the asset and from residual dependence on virgin-material purchases.	Structural comparison (4.2.1)	Tukker, 2004; Reim et al., 2015; Cimprich et al., 2023
CRM Price Volatility Exposure	The archetype's economic exposure to volatile virgin CRM prices, considering the extent to which future material needs can be covered by material returned under the archetype.	Context-derived (4.2.2)	Gaustad et al., 2018; Cimprich et al., 2023

Tab. 4.2 - Final decision variables in the Economic/Business Case macro-category.

4.3.3 Technical/Product

This category groups characteristics of the asset and the technology required to deliver, monitor and verify the offer. These are treated as properties of the archetype rather than of any one OEM's actual capacity to build them, which is addressed separately under Organisational/Capabilities.

Technological Infrastructure and Complexity Requirement records the telemetry, monitoring and service-delivery infrastructure that an archetype requires the OEM to provide (Sections 4.1.2, 4.2.1; Baines et al., 2007; Hidalgo-Crespo et al., 2025).

Product Reliability and Durability records the structural incentive an archetype creates for the OEM to design the asset for durability and reliable long-term performance. This incentive

is driven by whether the OEM retains an ongoing stake in the asset's condition, rather than by any technical outcome that the archetype itself can guarantee (Sections 4.1, 4.2.1; Tukker, 2015; Bocken et al., 2016; Reim et al., 2015).

Technology and Material Substitution Risk over the Contract Horizon records the risk that the embedded technology or materials become obsolete, or economically disadvantageous relative to available alternatives, before the archetype's commitment horizon elapses. This is distinct from Exposure to Regulatory Change over the Contract Horizon below, which concerns a change in the applicable rule rather than in what is technically or economically viable. This exposure resets at each transaction under a product-oriented offer but persists for the other two (Section 4.2.2; Tukker, 2015; Hidalgo-Crespo et al., 2025).

Decision variable	Description	Origin	Citations
Technological Infrastructure and Complexity Requirement	How technologically equipped this archetype requires the OEM to be, in monitoring and service-delivery infrastructure.	Structural comparison (4.2.1)	Baines et al., 2007; Hidalgo-Crespo et al., 2025
Product Reliability and Durability	The structural incentive this archetype creates for the OEM to design the asset for durability and reliable long-term performance, rather than a guaranteed technical outcome.	Structural comparison (4.2.1)	Tukker, 2015; Reim et al., 2015; Bocken et al., 2016
Technology and Material Substitution Risk over the Contract Horizon	The risk that the embedded technology or materials become obsolete, or economically disadvantageous relative to available alternatives, before the commitment horizon this archetype locks the OEM into has elapsed; distinct from a change in the applicable rule.	Context-derived (4.2.2)	Tukker, 2015; Hidalgo-Crespo et al., 2025

Tab. 4.3 - Final decision variables in the Technical/Product macro-category.

4.3.4 Organisational/Capabilities

This category groups what an archetype requires from the OEM's own organisation. Every variable in this category is defined at the level of the archetype rather than that of any one company: it records the requirement structurally imposed on whichever OEM adopts it, not the capability level a specific OEM currently holds.

Service and Maintenance Capability records the extent of operating capability an archetype structurally requires the OEM to have in order to deliver or coordinate maintenance during use. This is distinct from the coordination effort required to secure the asset's return, captured below. Even a product-oriented sale carries a baseline warranty obligation; what changes across archetypes is the extent to which this baseline is exceeded,

rising to the full, continuous operating responsibility assumed by a result-oriented offer (Tukker, 2004; Reim et al., 2015; Baines et al., 2007; Frederiksen et al., 2021). Despite its name, a higher rating denotes a greater requirement structurally imposed on the OEM, not a more favourable capability, and the variable is scored as a cost criterion accordingly.

Organisational Restructuring Requirement records the extent of internal transformation in roles, incentives and decision-making required by the use- or result-oriented CBM archetypes (Section 4.2.1; Baines et al., 2009; Frederiksen et al., 2021).

Decision variable	Description	Origin	Citations
Service and Maintenance Capability	How much this archetype requires the OEM to be able to deliver or coordinate maintenance during use, rising gradually from a baseline warranty obligation to full operating responsibility.	Structural comparison (4.2.1)	Tukker, 2004; Reim et al., 2015; Baines et al., 2007; Frederiksen et al., 2021
Organisational Restructuring Requirement	How much organisational transformation this CBM archetype's relational/result-based model requires of the OEM.	Structural comparison (4.2.1)	Baines et al., 2009; Frederiksen et al., 2021; Muto et al., 2016

Tab. 4.4 - Final decision variables in the Organisational/Capabilities macro-category.

4.3.5 Value Chain/Ecosystem

This category groups the OEM's relationships with actors outside its own organisational boundary, including reverse-logistics and recycling partners, third-party service providers, and technology and component subcontractors on which it depends but over which it has no direct control.

CRM Recovery and Asset Return Effectiveness records how effectively embedded critical material physically returns to, and is placed under the control of, the OEM at end of life (Sections 4.1, 4.2.1; Yang et al., 2018; Gaustad et al., 2018; Govindan et al., 2015). It does not extend to the technical feasibility of recycling the returned material once it is back under the OEM's control, which this thesis treats as a separate, technology-dependent parameter.

Lifecycle Data Accessibility and Control records governance over who controls and can act on the asset's own lifecycle data. Its scope is limited to the asset itself rather than to any wider system within which it may operate (Section 4.2.1; Reim et al., 2015; Yang et al., 2018). It is distinct from the technology required to generate that data, captured above.

Ecosystem Coordination Complexity records how complex the external network the OEM must coordinate becomes to sustain the offer and secure the asset's return. Its rating does not decline smoothly with ownership retention. The rating for the product-oriented archetype assumes that the OEM pursues the optional take-back mechanism described in Section 4.1.1, whose use depends on the customer's decision to return the product; in that case, the OEM

depends on reverse-logistics infrastructure it cannot fully control (Govindan et al., 2015). A use-oriented OEM works through an established, standardised partner network (Reim et al., 2015), and a result-oriented OEM internalises the core asset but depends on a wider constellation of specialised subcontractors (Hidalgo-Crespo et al., 2025; Muto et al., 2016). The coordination burden is therefore high at both ends of the spectrum, as established in Section 4.2.1. The rating describes this steady-state coordination burden, once each archetype’s ecosystem is already in place, rather than the transitional effort of assembling it.

Decision variable	Description	Origin	Citations
CRM Recovery and Asset Return Effectiveness	The extent to which the asset carrying the critical material physically returns to, and is placed under the operational control of, the OEM at end of life, independent of whether that material is technically recyclable once recovered.	Structural comparison (4.2.1)	Yang et al., 2018; Gaustad et al., 2018; Govindan et al., 2015; Pollard et al., 2023
Lifecycle Data Accessibility and Control	Governance over who controls, and can act on, the asset’s lifecycle data, distinct from the infrastructure needed to generate it.	Structural comparison (4.2.1)	Reim et al., 2015; Hidalgo-Crespo et al., 2025
Ecosystem Coordination Complexity	The complexity of the external network of actors the OEM must coordinate to sustain the offer during use and to secure the asset’s return and recovery under this archetype.	Structural comparison (4.2.1)	Govindan et al., 2015; Reim et al., 2015; Hidalgo-Crespo et al., 2025; Muto et al., 2016

Tab. 4.5 - Final decision variables in the Value Chain/Ecosystem macro-category.

4.3.6 Regulatory/Institutional

This category groups the legal and policy environment surrounding the OEM: obligations that apply regardless of the archetype, and incentives whose effects are not neutral across the three.

Regulatory Traceability Compliance Capability records the extent to which an archetype provides the OEM with structural conditions that support traceability and disclosure compliance, rather than the compliance capacity a specific OEM happens to hold. The relevant obligations include the labelling, data-carrier and recovered-material disclosure requirements that the EU Critical Raw Materials Act sets for permanent-magnet-containing products above a defined threshold (Regulation (EU) 2024/1252, Arts. 28–29), as well as the updatable, part-level product information that the Ecodesign for Sustainable Products Regulation’s digital product passport is designed to hold. As with Service and Maintenance Capability above, this is an archetype-level construct. This variable concerns the OEM’s structural position to meet compliance obligations, not who controls the underlying data. An

archetype under which the OEM retains continuous responsibility for the asset structurally places it in a position to maintain traceable lifecycle records, including maintenance, replacement and component history, that can support compliance with these and future disclosure requirements. These records should not, however, be equated with the specific dataset required by a digital product passport. The formal obligation applies to every archetype regardless of this variable's rating and does not by itself discriminate between them. The construct is therefore distinct from Lifecycle Data Accessibility and Control above: that variable concerns governance over the asset's data as such, whereas this one concerns whether the archetype structurally positions the OEM to use that data, together with the records generated through its continuing responsibility, to satisfy external disclosure obligations.

Regulatory and Financial Incentive Availability records whether fiscal or regulatory incentives supporting ownership-retaining models exist in the OEM's jurisdiction (Section 4.2.2; Assmann et al., 2023; Geissdoerfer et al., 2023), without assuming that a given OEM automatically qualifies for or claims them.

Exposure to Regulatory Change over the Contract Horizon records the archetype's exposure to changes in the regulatory framework during the OEM's period of commitment. This exposure resets at each transaction under a product-oriented offer but persists for the other two (Section 4.2.2; Cimprich et al., 2023; Tukker, 2015; Regulation (EU) 2024/1252).

Decision variable	Description	Origin	Citations
Regulatory Traceability Compliance Capability	The extent to which the archetype gives the OEM structural conditions supporting traceability and disclosure compliance on embedded critical materials, rather than a specific OEM's actual compliance capacity.	Context-derived (4.2.2)	Pollard et al., 2023; Regulation (EU) 2024/1252
Regulatory and Financial Incentive Availability	Whether regulatory or fiscal incentives exist that are designed to support ownership-retaining CBMs.	Context-derived (4.2.2)	Assmann et al., 2023; Geissdoerfer et al., 2023
Exposure to Regulatory Change over the Contract Horizon	The archetype's exposure to the regulatory framework changing during the period of commitment it locks the OEM into, distinct from changes in the technically or economically viable technology or material.	Context-derived (4.2.2)	Cimprich et al., 2023; Tukker, 2015; Regulation (EU) 2024/1252

Tab. 4.6 - Final decision variables in the Regulatory/Institutional macro-category.

Each of the eighteen variables above is traced to one of the two origins recorded in its table and kept analytically distinct from its closest neighbours. Section 4.4 now rates each of them,

for each archetype, on the ordinal scale and benefit/cost classification established in Chapter 3.

4.4 Rating the Archetypes and Testing Scenario Sensitivity

Each variable is assigned a Low/Medium/High rating and a benefit or cost classification for each archetype, following the ordinal scale and benefit/cost logic established in Chapter 3 (Sections 3.2.5-3.2.6), then converted into the numerical values the weighted scoring procedure requires: 0.33, 0.67, 1.00. Together they form the baseline CBM performance matrix, a reference configuration representing a neutral, intermediate CRM market condition rather than a claim about current market conditions; the Low and High deviations tested below are ordinal states defined relative to this baseline, not independent market forecasts. Six of the eighteen variables, those whose mechanisms link directly to CRM price and predictability, are then tested for their response when that market is allowed to move away from baseline; the remaining twelve are unaffected and are not repeated.

4.4.1 Assigning Low/Medium/High Ratings and Benefit/Cost Classification across the Three Archetypes

Each rating below is a relative judgement among the three archetypes for that variable, not an absolute measure: a Low rating for capital requirement means only that the product-oriented archetype ties up comparatively little capital, not that it ties up none. Each rating follows directly from the mechanism established for that variable in Sections 4.2 and 4.3 and is not reargued here.

Market. Customer Acceptance of the CBM's Ownership/Access Arrangement (benefit) is High/Medium/Low for product/use/result, since resistance to an unfamiliar arrangement increases as the offer moves from a conventional sale towards an access-based arrangement (Assmann et al., 2023; Tukker, 2015). Customer Asset-Care and Moral-Hazard Exposure (cost) is Low/High/Medium: exposure is lowest under the product-oriented configuration, where the customer owns the asset outright; highest under the use-oriented configuration, where the customer operates an asset it does not own under comparatively limited OEM oversight; and lower, though not eliminated, under the result-oriented configuration. Here, the OEM's retained decision authority over the core asset's technical means and maintenance limits the customer's opportunity to under-care for it, while the customer's residual responsibilities for the wider system in which the asset operates still carry a documented, if smaller, non-ownership care risk (Tukker, 2015; Reim et al., 2015). Demand/Usage Forecasting Dependency (cost) is Low/Medium/High: use's revenue depends on the arrangement being renewed and access continuing to be taken up over the contract, while result's economics depend on forecasting the volume or intensity of demand accurately enough to price and resource a fixed, pre-agreed commitment, so that any forecasting error is absorbed entirely by the OEM rather than passed through in revenue, a more exacting and harder-to-forecast basis that places it above use (Reim et al., 2015; Assmann et al., 2023).

Decision variable	Benefit/Cost	Product	Use	Result	Citations
Customer Acceptance of the CBM's Ownership/Access Arrangement	Benefit	High	Medium	Low	Tukker, 2015; Assmann et al., 2023
Customer Asset-Care and Moral-Hazard Exposure	Cost	Low	High	Medium	Tukker, 2015; Reim et al., 2015
Demand/Usage Forecasting Dependency	Cost	Low	Medium	High	Reim et al., 2015; Assmann et al., 2023

Tab. 4.7 - Ratings for the Market macro-category.

Economic/Business Case. Provider Investment/Capital Requirement (cost) is Low/Medium/High, following the capital gradient established in Section 4.2.1 (Tukker, 2004; Reim et al., 2015; Tukker, 2015). Operational Service-Delivery Cost Burden (cost) is Low/Medium/High, following the increase in operational responsibility as the offer moves towards a sustained relationship (Reim et al., 2015; Tukker, 2015). Value Creation and Capture Potential (benefit) is Medium/High/High: ownership retention gives use and result the property-rights position from which value generated over the relationship's life, including the value of material recovered at end of life, can be routed back into the OEM's own chain (Section 4.3.2), and the two are rated equally because the greater asset control under result is offset by the capture risk associated with an outcome that may not always be straightforward to define or measure (Tukker, 2004; Cimprich et al., 2023). CRM Price Volatility Exposure (cost) is High/Medium/Low, following the extent to which recovered material can meet each archetype's material needs (Gaustad et al., 2018; Cimprich et al., 2023).

Decision variable	Benefit/Cost	Product	Use	Result	Citations
Provider Investment/Capital Requirement	Cost	Low	Medium	High	Tukker, 2004; Reim et al., 2015; Tukker, 2015
Operational Service-Delivery Cost Burden	Cost	Low	Medium	High	Reim et al., 2015; Tukker, 2015
Value Creation and Capture Potential	Benefit	Medium	High	High	Tukker, 2004; Cimprich et al., 2023; Yang et al., 2018
CRM Price Volatility Exposure	Cost	High	Medium	Low	Gaustad et al., 2018; Cimprich et al., 2023

Tab. 4.8 - Ratings for the Economic/Business Case macro-category.

Technical/Product. Technological Infrastructure and Complexity Requirement (cost) is Low/Medium/High, since use- and result-oriented offers require sensing and data-platform capabilities that a sale does not (Baines et al., 2007; Muto et al., 2016; Hidalgo-Crespo et al., 2025). Product Reliability and Durability (benefit) is Low/High/High: use and result are rated equally because both create the same ownership-retention incentive to design for durability and reliable long-term performance (Tukker, 2015; Reim et al., 2015; Bocken et al., 2016). Technology and Material Substitution Risk over the Contract Horizon (cost) is Low/High/Medium. Product resets this exposure at each transaction and is therefore rated Low. Use is rated High because the contract ties the OEM, and the customer, to the specific asset and technical solution installed at signing. Substituting the underlying technology during the term is not precluded outright but is more constrained by the contract and the configuration of the installed asset than under the other two archetypes. Result is rated Medium because the OEM's contractual commitment is to a functional result rather than to a specific technical means, preserving the design freedom established in Section 4.2.1 to adapt the underlying technology during the contract, even though its commitment is no shorter-lived than that of use (Tukker, 2015; Hidalgo-Crespo et al., 2025).

Decision variable	Benefit/Cost	Product	Use	Result	Citations
Technological Infrastructure and Complexity Requirement	Cost	Low	Medium	High	Baines et al., 2007; Muto et al., 2016; Hidalgo-Crespo et al., 2025
Product Reliability and Durability	Benefit	Low	High	High	Tukker, 2015; Reim et al., 2015; Bocken et al., 2016; Yang et al., 2018
Technology and Material Substitution Risk over the Contract Horizon	Cost	Low	High	Medium	Tukker, 2015; Hidalgo-Crespo et al., 2025

Tab. 4.9 - Ratings for the Technical/Product macro-category.

Organisational/Capabilities. Both variables are rated at the archetype level rather than at the company level (Section 4.3.4). Service and Maintenance Capability (cost) is Low/Medium/High, reflecting a requirement that ranges from a baseline warranty obligation to full operating responsibility (Tukker, 2004; Reim et al., 2015; Baines et al., 2007). Organisational Restructuring Requirement (cost) is Low/Medium/High, reflecting progressively greater internal transformation from product- to result-oriented offers (Baines et al., 2009; Frederiksen et al., 2021; Muto et al., 2016).

Decision variable	Benefit/Cost	Product	Use	Result	Citations
Service and Maintenance Capability	Cost	Low	Medium	High	Tukker, 2004; Reim et al., 2015; Baines et al., 2007
Organisational Restructuring Requirement	Cost	Low	Medium	High	Baines et al., 2009; Frederiksen et al., 2021; Muto et al., 2016

Tab. 4.10 - Ratings for the Organisational/Capabilities macro-category.

Value Chain/Ecosystem. CRM Recovery and Asset Return Effectiveness (benefit) is Low/Medium/High: use is rated below result because, although hand-back is structurally incorporated into the lease arrangement at each contract term (Section 4.1.2), it still constitutes a discrete return event whose timing and condition the OEM does not fully control until it occurs. Under result, by contrast, the OEM retains control throughout and no separate hand-back event is required (Yang et al., 2018; Gaustad et al., 2018). Lifecycle Data Accessibility and Control (benefit) is Low/Medium/High, following the visibility gradient established in Section 4.2.1 (Reim et al., 2015; Yang et al., 2018). Ecosystem Coordination Complexity (cost) is High/Medium/High, reproducing the non-monotonic pattern established in Section 4.2.1: product's rating assumes the OEM pursues a take-back mechanism (Section 4.1), whose use still depends on the customer's decision to return the product; without such a mechanism the archetype would face no equivalent coordination burden (Govindan et al., 2015; Reim et al., 2015; Hidalgo-Crespo et al., 2025; Muto et al., 2016). CRM Recovery and Ecosystem Coordination Complexity capture distinct aspects of the same mechanism: the first records the outcome, namely how much material actually returns under the OEM's control, while the second records the burden of coordinating the actors involved. For product, the first is low because return depends on the customer's decision, and the second is high because the OEM depends on reverse-logistics and recycling actors it does not control.

Decision variable	Benefit/Cost	Product	Use	Result	Citations
CRM Recovery and Asset Return Effectiveness	Benefit	Low	Medium	High	Yang et al., 2018; Gaustad et al., 2018; Govindan et al., 2015
Lifecycle Data Accessibility and Control	Benefit	Low	Medium	High	Reim et al., 2015; Yang et al., 2018
Ecosystem Coordination Complexity	Cost	High	Medium	High	Govindan et al., 2015; Reim et al., 2015; Hidalgo-Crespo et al., 2025; Muto et al., 2016

Tab. 4.11 - Ratings for the Value Chain/Ecosystem macro-category.

Regulatory/Institutional. Regulatory Traceability Compliance Capability (benefit) is Low/Medium/High. Product is rated Low because the archetype leaves the OEM with more limited direct lifecycle visibility after sale, increasing its dependence on customer or third-party information to reconstruct any subsequent record. Use and result both structurally give the OEM continuing responsibility for the asset and, with it, the lifecycle, maintenance and component records generated through that responsibility. The two differ, however, in how central this record-keeping is to the offer itself: a use-oriented OEM tracks the asset primarily to support access and maintenance, whereas a result-oriented OEM's continuing responsibility for contracted performance makes comprehensive, ongoing record-keeping integral to delivering and verifying that performance. This difference justifies the higher rating for result (Pollard et al., 2023; Regulation (EU) 2024/1252). Regulatory and Financial Incentive Availability (benefit) is Low/High/High: use and result are rated equally because the evidence supports an asymmetry between product-oriented and ownership-retaining archetypes rather than a further gradation between the two (Assmann et al., 2023; Geissdoerfer et al., 2023). Exposure to Regulatory Change over the Contract Horizon (cost) is Low/High/High, since neither ownership-retaining archetype provides a reset point within a single agreement (Cimprich et al., 2023; Tukker, 2015; Regulation (EU) 2024/1252; Ragonnaud, 2024).

Decision variable	Benefit/Cost	Product	Use	Result	Citations
Regulatory Traceability Compliance Capability	Benefit	Low	Medium	High	Pollard et al., 2023; Regulation (EU) 2024/1252
Regulatory and Financial Incentive Availability	Benefit	Low	High	High	Assmann et al., 2023; Geissdoerfer et al., 2023
Exposure to Regulatory Change over the Contract Horizon	Cost	Low	High	High	Tukker, 2015; Regulation (EU) 2024/1252; Ragonnaud, 2024

Tab. 4.12 - Ratings for the Regulatory/Institutional macro-category.

No archetype dominates across all eighteen variables: product rates favourably on capital exposure and customer acceptance but poorly on recovery and regulatory incentive availability; result shows the reverse profile, combined with the highest organisational and infrastructural demands; use occupies an intermediate position overall, though not on every variable. Resolving this trade-off is the purpose of the multi-criteria structure, once these ratings are normalised and weighted through the AHP.

4.4.2 Normalising the Ratings into the Baseline CBM Performance Matrix

Section 3.3.4 converts each Low/Medium/High rating into a numerical score in three steps. First, the rating is encoded numerically irrespective of benefit/cost classification, with Low = 1, Medium = 2 and High = 3 (Equation 3.14). Second, the scores are harmonised by benefit/cost classification so that a higher value always represents a more favourable

outcome for the archetype: the score remains equal to the raw encoding for a benefit variable, whereas for a cost variable it is calculated as four minus the raw encoding (Equation 3.15; Hwang & Yoon, 1981). Third, the harmonised score is normalised by dividing it by three, yielding 0.33, 0.67 or 1.00 (Equation 3.16). Applying this procedure to the eighteen ratings in Section 4.4.1 produces the baseline CBM performance matrix, P, formally defined in Section 3.3.4. Harmonisation means that the same raw High rating can normalise to 1.00 for a benefit variable (Customer Acceptance, product) or to 0.33 for a cost variable (CRM Price Volatility Exposure, product). A variable's numerical score therefore cannot be inferred from its rating alone without first checking its benefit/cost classification.

Macro-category	Decision variable	Benefit/Cost	Product	Use	Result
Market	Customer Acceptance of the CBM's Ownership/Access Arrangement	Benefit	1.00	0.67	0.33
Market	Customer Asset-Care and Moral-Hazard Exposure	Cost	1.00	0.33	0.67
Market	Demand/Usage Forecasting Dependency	Cost	1.00	0.67	0.33
Economic/Business Case	Provider Investment/Capital Requirement	Cost	1.00	0.67	0.33
Economic/Business Case	Operational Service-Delivery Cost Burden	Cost	1.00	0.67	0.33
Economic/Business Case	Value Creation and Capture Potential	Benefit	0.67	1.00	1.00
Economic/Business Case	CRM Price Volatility Exposure	Cost	0.33	0.67	1.00
Technical/Product	Technological Infrastructure and Complexity Requirement	Cost	1.00	0.67	0.33
Technical/Product	Product Reliability and Durability	Benefit	0.33	1.00	1.00
Technical/Product	Technology and Material Substitution Risk over the Contract Horizon	Cost	1.00	0.33	0.67

Macro-category	Decision variable	Benefit/Cost	Product	Use	Result
Organisational/Capabilities	Service and Maintenance Capability	Cost	1.00	0.67	0.33
Organisational/Capabilities	Organisational Restructuring Requirement	Cost	1.00	0.67	0.33
Value Chain/Ecosystem	CRM Recovery and Asset Return Effectiveness	Benefit	0.33	0.67	1.00
Value Chain/Ecosystem	Lifecycle Data Accessibility and Control	Benefit	0.33	0.67	1.00
Value Chain/Ecosystem	Ecosystem Coordination Complexity	Cost	0.33	0.67	0.33
Regulatory/Institutional	Regulatory Traceability Compliance Capability	Benefit	0.33	0.67	1.00
Regulatory/Institutional	Regulatory and Financial Incentive Availability	Benefit	0.33	1.00	1.00
Regulatory/Institutional	Exposure to Regulatory Change over the Contract Horizon	Cost	1.00	0.33	0.33

Tab. 4.13 - Baseline CBM performance matrix: normalised scores (0.33 / 0.67 / 1.00) for all eighteen decision variables across the three archetypes.

One consequence of this harmonisation is worth noting because it reflects a structural pattern rather than a numerical artefact. For Ecosystem Coordination Complexity, product and result both score 0.33 despite being rated High, the least favourable level, but for structurally opposite reasons: a product-oriented OEM depends on reverse-logistics infrastructure it cannot fully control, whereas a result-oriented OEM internalises the core asset but depends on a wider constellation of specialised subcontractors (Section 4.4.1). This numerical pattern mirrors the qualitative one: the two archetypes at opposite ends of the ownership spectrum can reach comparably burdensome outcomes through different structural routes, while the intermediate use-oriented archetype, scoring 0.67, does not necessarily do so. This matrix, together with the AHP-derived weights, provides the input to the scoring procedure in Section 3.3.5 (Equations 3.17-3.18).

4.4.3 Testing the Baseline against Low and High Critical-Raw-Material Market Scenarios

As Section 3.4.1 sets out, a chosen archetype does not operate independently of the CRM market on which it depends. The scenario logic treats Low, Baseline and High conditions as joint states of price level and predictability by design choice, not as a claim that the two necessarily co-move for critical raw materials (Section 3.4.1). Of the eighteen variables, six respond directly to this joint condition, each reflecting a distinct aspect of how that condition affects the CBM archetype: CRM Price Volatility Exposure (dependence on virgin-material procurement); Value Creation and Capture Potential (value of recoverable material); Technology and Material Substitution Risk over the Contract Horizon (substitution risk); Operational Service-Delivery Cost Burden (labour cost of repair versus replacement); Service and Maintenance Capability (organisational requirement for that same repair-intensive maintenance); and Customer Acceptance of the CBM's Ownership/Access Arrangement (customer sensitivity to the price and financing risk of outright ownership). These six are the variables for which the scenario indicator c_i of Equation 3.20 takes the value 1; for the remaining twelve, $c_i = 0$. A variable is treated as scenario-sensitive only where the literature reviewed provides either direct empirical evidence or a theoretically grounded causal mechanism linking it to CRM price or predictability. The remaining twelve variables have no such documented mechanism or record a structural feature that does not vary with it and retain their baseline rating under every scenario. They are therefore not repeated here. Each variable below is rated Low/Medium/High under the Low and High conditions, following its established benefit/cost classification, and normalised exactly as in Section 4.4.2. Every movement is bounded to at most one step on the three-level scale in either direction from baseline. Scenario ratings are expressed on the same Low/Medium/High scale as the baseline and record how far each rating moves from its baseline value. Where a rating is already at the scale's ceiling or floor, it cannot move further in that direction; this is a limitation of the three-level instrument rather than evidence that the underlying exposure stops changing.

CRM Price Volatility Exposure maps onto the closed-loop/open-loop distinction: an OEM retaining ownership keeps material stock within its own boundary, creating a closed-loop condition insulated from the open market, whereas an OEM that sells outright loses that stock to an open-loop flow (Cimprich et al., 2023). Under a Low market, product improves by one step, from High to Medium; the bound of one step prevents it from reaching Low. Use converges to result's already favourable Low, since a calmer market narrows the practical gap between partial and negligible dependence. Under a High market, product is already at its ceiling and cannot deteriorate further, while use and result retain their baseline Medium and Low. This follows from a modelling assumption, discussed after Tab. 4.14: material returned to the OEM is assumed to keep the exposure of the ownership-retaining archetypes to virgin-material prices below that of product.

Value Creation and Capture Potential. This variable tracks the market value the OEM can capture by monetising recovered material as CRM prices shift; the archetype's avoided dependence on virgin-material purchases is captured separately, and through a different mechanism, under CRM Price Volatility Exposure above, so the same price movement is not counted twice by the two variables. Under a Low market, the market value of recoverable material declines, so use and result's baseline advantage over product largely disappears, with both falling to product's unchanged Medium. Under a High market, the same material

becomes sufficiently valuable that use and result, which retain it, would gain the most, but their rating is already at the ceiling and cannot improve further. The widening of the gap is therefore recorded as a fall of product, which sold the asset and forgoes that value, to Low. This is a movement relative to use and result, not a loss of value that product previously held.

Technology and Material Substitution Risk over the Contract Horizon. The future market share of rare-earth-based designs depends closely on how rare-earth prices move relative to substitutes, with rising prices and supply disruption being principal drivers of substitution (Pavel et al., 2017). Product, which resets its exposure at each transaction, remains at Low under Low and Baseline conditions but rises to Medium under a High market, since even a single-transaction sale must be designed within an accelerating substitution environment. Use, whose contract ties the OEM to the specific technology installed at signing for the duration of the lease, is rated at its baseline High. Under a Low market, that constraint matters less because a calmer substitution environment gives the installed technology more time before it is overtaken, easing the rating to Medium. Under a High market, it remains at High, its ceiling. Result, whose functional-result commitment preserves the freedom to adapt the underlying technology during the contract, remains less exposed than use precisely because of that freedom, but the freedom to adapt does not eliminate the underlying pressure to do so: at its baseline Medium, a Low market eases the rating further to Low, since little adaptation is actually required, while a High market raises it to High, since the same acute price and scarcity pressure that drives substitution generally also raises the cost of redesign, requalification and integration that adapting the technical solution requires, so result rises to High, reaching the same rating as use under a High market, although its contractual freedom to adapt provides greater flexibility. This is a risk of technical and economic substitution becoming disadvantageous, not a direct exposure to the OEM's own virgin-material purchase price, which is captured separately under CRM Price Volatility Exposure above.

Operational Service-Delivery Cost Burden and Service and Maintenance Capability respond to the same underlying market shift, but along two distinct dimensions rather than as a repeated effect: under CRM price pressure, OEMs responsible for maintenance shift labour from replacing worn components towards repairing and reconditioning them, reflecting a documented trade-off between labour costs and the cost of capital and new parts (Cimprich et al., 2023; Gaustad et al., 2018; Lüdeke-Freund et al., 2019). Cost Burden records the financial consequence of that shift; Service and Maintenance Capability records its organisational consequence. For Operational Service-Delivery Cost Burden, this affects use and result directly because both carry direct maintenance responsibility. Result improves to Medium under a Low market but remains at High, its ceiling, under a High market. Use moves across the full scale, from Low under a Low market to High under a High market. Product is only partly exposed because its transaction is settled once, but even its baseline warranty becomes harder to discharge as replacement components become scarcer. Its rating therefore rises to Medium under a High market while remaining at Low otherwise.

Service and Maintenance Capability follows the same logic at the organisational rather than cost level. Under a Low market, no rating changes, because organisational requirements are rigid downward: a calmer market does not reduce the operating structure already

required. Under a High market, use worsens to High because it lacks result's already-integrated operating structure, while product worsens to Medium for the same reason as its cost-burden counterpart. Result, already at High, cannot deteriorate further.

Customer Acceptance of the CBM's Ownership/Access Arrangement. Capital-goods leasing research documents a shift in customer preference towards access-based arrangements as the price and financing risk of ownership rise, since leasing transfers that risk to the OEM (Frederiksen et al., 2021; Reim et al., 2015; Tukker, 2004). Baseline ratings hold under a Low market because a calm market provides no additional reason to reconsider ownership. Under a High market, product falls to Medium, use improves to High, and result improves to Medium, since acute scarcity can also increase customers' willingness to transfer performance risk entirely to the OEM (Cimprich et al., 2023; Gaustad et al., 2018).

Decision variable	Benefit/Cost	Archetype	Low	Baseline	High
CRM Price Volatility Exposure	Cost	Product	0.67	0.33	0.33
		Use	1.00	0.67	0.67
		Result	1.00	1.00	1.00
Value Creation and Capture Potential	Benefit	Product	0.67	0.67	0.33
		Use	0.67	1.00	1.00
		Result	0.67	1.00	1.00
Technology and Material Substitution Risk over the Contract Horizon	Cost	Product	1.00	1.00	0.67
		Use	0.67	0.33	0.33
		Result	1.00	0.67	0.33
Operational Service-Delivery Cost Burden	Cost	Product	1.00	1.00	0.67
		Use	1.00	0.67	0.33
		Result	0.67	0.33	0.33
Service and Maintenance Capability	Cost	Product	1.00	1.00	0.67
		Use	0.67	0.67	0.33
		Result	0.33	0.33	0.33
Customer Acceptance of the CBM's Ownership/Access Arrangement	Benefit	Product	1.00	1.00	0.67
		Use	0.67	0.67	1.00
		Result	0.33	0.33	0.67

Tab. 4.14 - Scenario-dependent decision variables: normalised scores under the Low, Baseline and High critical-raw-material market conditions.

Tab. 4.14 also shows that the result-oriented archetype's score moves on only two of the six variables between the baseline and high-market columns: Customer Acceptance and Technology and Material Substitution Risk over the Contract Horizon. This pattern is not imposed by the scenario test: on three of the other four variables it follows from where result already sits on the scale, and on the fourth from the modelling assumption noted below.

Result does not move on the other four variables. On three of them, Value Creation and Capture Potential, Operational Service-Delivery Cost Burden and Service and Maintenance Capability, it already sits at the ceiling or floor of the three-level scale for reasons established independently in Section 4.4.1, on grounds unconnected to the severity of any future CRM market: its full property-rights position over value created, and the organisational and cost commitment its continuous operating responsibility already assumes at baseline. A High market cannot push an already-extreme rating further in the same direction, the ceiling-and-floor limitation of the three-level instrument noted above; this reflects a limit of the measurement scale, not evidence that result's underlying exposure is unaffected by market conditions. For CRM Price Volatility Exposure, the unchanged ratings of use and result under the High market are a modelling assumption rather than a limit of the scale. Because these archetypes retain greater ownership or control over the asset and its materials, the model assumes that the potential recovery of material at end of life reduces their structural exposure to virgin-material prices relative to product. This assumption is favourable to use and result. For a durable asset, however, the material available for recovery becomes largely available only at end of life, meaning that the corresponding benefit may arise only late in the contract horizon. Technology and Material Substitution Risk and Customer Acceptance are, by contrast, the two variables where result does not start at the extreme in the direction of the high-market shift: Medium for Technology and Material Substitution Risk over the Contract Horizon, which can rise to High, and Low for Customer Acceptance, which can rise to Medium. There is therefore room for the archetype's design freedom and property-rights position to be tested, rather than only reconfirmed, once the underlying market condition moves.

For the twelve variables not listed, the matrix is unchanged across all three market conditions; for the six above, the Baseline column reproduces Tab. 4.13 for comparison. Together with the AHP weights, Tab. 4.14 completes the input to the scenario scoring procedure (Equations 3.19-3.21, Section 3.4.2), producing a weighted score for each archetype under each market condition.

4.4.4 Trade-Off Profile of the Three Archetypes across Macro-Categories

The ratings of Sections 4.4.1 to 4.4.3 can be summarised at the level of the six macro-categories. Fig. 4.1 reports, for each archetype under the baseline CRM market condition, the simple average of the normalised variable scores of Tab. 4.13 within each category, on a 0-1 scale where higher is more favourable. No importance weights are applied at this stage, neither to the variables within a category nor to the categories themselves: the figure is therefore descriptive and is not part of the DSS scoring, which combines these scores with the weights derived in Chapter 5.




Fig. 4.1 - Structural profile of the three CBM archetypes under the baseline CRM market condition.

The figure makes visible the trade-off noted at the end of Section 4.4.1. The product-oriented archetype is the most favourable on Market and Organisational/Capabilities (both 1.00) and on Technical/Product (0.78), but the least favourable on Value Chain/Ecosystem (0.33) and Regulatory/Institutional (0.55). The result-oriented archetype shows the opposite profile, with its highest values on Value Chain/Ecosystem and Regulatory/Institutional (both 0.78) and its lowest on Organisational/Capabilities (0.33) and Market (0.44). The use-oriented archetype occupies an intermediate position on most categories. Economic/Business Case is the least differentiated category (0.75, 0.75 and 0.67 for product, use and result). Because no archetype is favourable across all categories, which one is preferable depends on the relative importance attributed to the categories, which is the role of the AHP weighting.

The scenario-specific ratings of Tab. 4.14 are not represented here; their effect on the profile of each archetype is examined in Chapter 5.

4.5 Decision Support System Framework

Sections 4.1 to 4.4 progressively assembled the elements of the decision support system (DSS) developed in this thesis: the value chains through which the three CBM archetypes were mapped, the eighteen decision variables consolidated under six macro-categories, the Low, Medium and High ratings assigned to each archetype, and the resulting CBM performance matrix. None of these elements, taken individually, constitutes the DSS. The system becomes operational only once the importance judgements elicited from the experts consulted for this thesis, through the AHP, are combined with these literature-derived ratings and re-applied under varying CRM market conditions. This section describes that integration.

The resulting DSS combines four elements: AHP-based criteria weighting, an ordinal rating-based assessment of the three CBM archetypes, a weighted additive aggregation rule, and a scenario module that re-applies the performance matrix under alternative CRM market conditions while keeping the weights fixed. Section 4.5.1 presents the overall functional

architecture of the DSS, tracing the flow of information from the decision context established in Chapter 3 to the final recommendation. Section 4.5.2 details the two-level hierarchy through which the AHP derives the importance weights of the six macro-categories and eighteen decision variables. Section 4.5.3 explains how the three archetypes are scored against this weighted structure. Section 4.5.4 describes how the same architecture is re-applied when Low and High CRM market conditions replace the baseline configuration.

4.5.1 DSS Functional Architecture and Information Flow

DSSs have traditionally been described through architectures conceived at a level of generality that precedes any specific application. A widely used characterisation, for instance, distinguishes the language, presentation, knowledge and problem-processing subsystems through which a DSS can be described, regardless of the domain to which it is applied; different configurations of these subsystems give rise to different types of system (Holsapple, 2008).

The DSS developed in this thesis does not apply such a generic architecture and populate it with content from Chapters 3 and 4. Instead, it provides a functional representation of the specific decision problem addressed in this thesis. Each block corresponds to a construct or procedure established in the preceding sections, rather than to a component required by a general-purpose taxonomy. This is consistent with a broader premise underlying DSSs since they were first framed as a distinct class of information system: their purpose is to support the judgement of a specific decision-maker facing a specific class of semi-structured problem, rather than to conform to a predetermined architectural template (Sprague, 1980).

Fig. 4.2 should therefore be read as a problem-specific flowchart for the choice addressed in this thesis. It traces the constructs that make up the architecture, the sequence in which they are executed, the points at which the procedure loops back until a stated condition is satisfied, and the point at which it terminates.




Fig. 4.2 - DSS flowchart: from the decision problem and decision context to the baseline and scenario-specific DSS output.

Fig. 4.2 opens with the decision problem itself: which of the three CBM archetypes, mapped independently in Section 4.1, is most attractive for an OEM whose products embed CRMs, and for which the recovery of those materials is one consideration alongside cost, capital, market, organisational and regulatory conditions, under the baseline CRM market condition and, subsequently, under the Low and High conditions defined relative to it. This problem resolves, immediately beneath it in the flowchart, into the decision context established in Section 3.2.1: the OEM as decision-maker, the three archetypes as alternatives, and the six macro-categories fixed in Section 3.2.4, organising the eighteen decision variables consolidated in Section 4.3, as the criteria structure against which they are evaluated.

From this decision context, Fig. 4.2 divides into two input streams that the architecture keeps deliberately separate. The literature-based stream, shown in Fig. 4.2 as Literature-based inputs and grouped under Performance Assessment, supplies the Low, Medium or High ratings and the benefit/cost classification developed for each archetype in Section 4.4.1, reworked through the harmonisation and normalisation procedure of Sections 3.3.4 and 4.4.2 into the CBM performance matrix P of Tab. 4.13. This branch carries no dependency on expert judgement and remains identical across repeated applications of the framework, except for the six scenario-sensitive variables identified in Section 4.4.3, whose treatment is reserved for Section 4.5.4. The decision-makers' stream, shown in Fig. 4.2 as Decision-makers' inputs and grouped under AHP Weighting, is elicited directly from the experts consulted for this thesis, who act as the OEM's decision-makers, and consists exclusively of the pairwise comparisons through which the AHP derives the importance of each macro-category (Level 1) and of each decision variable within its category (Level 2), following the elicitation, consistency-checking and aggregation procedure of Sections 3.3.1 to 3.3.3. No qualitative rating in the performance matrix is elicited from the experts: their contribution is confined to the relative importance of the criteria, not to how the archetypes perform against them, which is what allows the AHP to be characterised precisely as a weighting mechanism. Within the AHP Weighting block, each of the two levels passes through two consistency checks, both drawn in Fig. 4.2 against the same threshold, $CR \leq 0.10$. The first is applied to the individual matrices of each expert: a matrix that fails it is returned to the expert for revision before any aggregation takes place. The matrices that pass are then aggregated across experts through the geometric mean (Equation 3.9), and the second check is applied to the resulting group matrix; should this fail, the procedure again returns to the experts' judgements at the corresponding level rather than to the aggregation step, since re-aggregating unchanged judgements would reproduce the same group matrix. Where the DSS is applied with a single respondent, the aggregated matrix coincides with the individual one and the second check with the first. Only once both checks are satisfied are the category weights (Level 1) and the local weights (Level 2) combined into the global weight W_i of each decision variable (Equation 3.10).

It is worth making explicit what this weighting module does not do. The pairwise comparisons are confined to the six macro-categories and the eighteen decision variables; the three archetypes are never placed in a pairwise comparison matrix and scored against one another through the same eigenvector procedure used to derive the category and variable weights. Extending the hierarchy to a third level of this kind would have implemented what has been termed relative measurement within the AHP (Saaty, 1986), adding fifty-four further comparisons per expert and making the experts' contribution depend on how the archetypes perform rather than only on how much each criterion matters. The architecture

instead assigns each archetype a rating on the fixed ordinal scale established in Section 3.3.4 and detailed for this case in Section 4.5.3: criterion importance is derived through comparison, while the alternatives are assessed through rating.

The two branches converge once the global weights W_i and the performance matrix P are both available. Every archetype's overall score is the weighted sum of its performance scores across the eighteen decision variables (Equation 3.17), following the additive logic through which criterion weights and performance scores are combined in multi-criteria decision analysis (Kaliszewski & Podkopaev, 2016); the recommended archetype under the baseline configuration is the one that maximises this score (Equation 3.18). In Fig. 4.2, this step corresponds to the Weighted aggregation block, whose result is the CBM ranking and recommended CBM for the current market condition, which on the first pass is the baseline configuration.

The architecture does not treat this baseline recommendation as the end point. Fig. 4.2 closes the loop through the Scenario selection check, which asks whether a market condition remains to be evaluated: while a scenario remains, the flowchart updates the performance matrix P only for the six scenario-sensitive variables identified in Section 4.4.3 leaves the remaining twelve variables and every global weight unchanged, and returns to the weighted-aggregation step (Equation 3.19) to identify the scenario-specific recommended archetype (Equation 3.21). Once both the Low and High CRM market conditions defined in Section 3.4.1 have been passed through this loop, no scenario remains and the flowchart terminates in the DSS output. This loop is deliberately distinct from the consistency-check loops described above: the latter are corrective, applied until a judgement satisfies a fixed threshold, whereas the scenario loop reruns an already specified model under a different, exogenously specified market condition, of the kind the scenario-planning literature treats as a predetermined element external to the choice under evaluation, whose consequences remain worth tracing systematically even though the element itself lies outside anyone's control (Schoemaker, 1995). Section 4.5.4 develops the substantive content of this mechanism.

Each pass through this loop, the baseline configuration and each of the two scenario deviations in turn, produces a CBM ranking and a recommended CBM, explicitly qualified in Fig. 4.2 as holding for the current market condition, that is, the one evaluated in that pass, so that the three outputs are never read as recommendations for the same market state. The DSS output, reached once no scenario remains, therefore consists of these three pass-specific results.

Both the baseline and the scenario-specific outputs are presented to the OEM as a recommendation to be interpreted, not a determination to be executed automatically, consistent with the distinction between decision-support tools that augment managerial judgement and systems designed to substitute for it, already drawn in Section 2.4.3. The OEM remains free to weigh the recommendation against considerations the DSS does not formalise.

4.5.2 AHP Weighting Hierarchy

This section applies the AHP elicitation, consistency-checking and aggregation procedure to its concrete object: the six macro-categories and eighteen decision variables consolidated in Section 4.3, which together form the two-level hierarchy against which the weights are derived.

The combinatorial argument formalised symbolically in Equations 3.3 and 3.4 of Section 3.3.1 can now be evaluated for the specific hierarchy developed here. Comparing the eighteen decision variables directly would require each expert to complete $n(n-1)/2 = 153$ pairwise judgements. Decomposing the same set under six macro-categories instead requires fifteen comparisons at Level 1 ($6 \times 5/2$) and a further nineteen distributed across the six category-level exercises at Level 2, following directly from the category sizes established in Section 4.3, namely three, four, three, two, three and three variables ($3 + 6 + 3 + 1 + 3 + 3 = 19$). The hierarchical decomposition adopted here therefore reduces the number of judgements an expert must supply from one hundred and fifty-three to thirty-four, a reduction of close to 4.5 times (Tab. 4.15).

Comparison set	Elements (n)	Comparisons $n(n-1)/2$
Direct comparison of all eighteen decision variables (no hierarchy)	18	153
Level 1 - six macro-categories	6	15
Level 2 - Market	3	3
Level 2 - Economic/Business Case	4	6
Level 2 - Technical/Product	3	3
Level 2 - Organisational/Capabilities	2	1
Level 2 - Value Chain/Ecosystem	3	3
Level 2 - Regulatory/Institutional	3	3
Level 2 subtotal	18	19
Hierarchical total required (Level 1 + Level 2)	–	34

Tab. 4.15 - Number of pairwise comparisons required: direct comparison of the eighteen decision variables versus the two-level hierarchical decomposition used to elicit expert judgements.

Fig. 4.3 sets out the resulting hierarchy in full. The global objective, the selection of the preferred archetype, sits at its head; the six macro-categories defined in Section 3.2.4 and populated in Section 4.3 occupy Level 1, each compared against the other five; the decision variables belonging to each category occupy Level 2, compared only against the other variables of the same category. The full definition of each variable, together with the rationale for its inclusion, was already given in Section 4.3 (Tabs. 4.1-4.6) and is not repeated here: Fig. 4.3 reproduces the variables by name alone, as the leaves of the structure against which the weighting procedure of Section 4.5.1 is applied.




Fig. 4.3 - AHP weighting hierarchy: global objective, six macro-categories and eighteen decision variables (variable definitions in Section 4.3).

Once the category weights W_c and the local weights within each category have been derived from this hierarchy, through the procedure detailed in Section 3.3.2, they are combined into the global weight of each decision variable by Equation 3.10. At that point, the AHP Weighting block of Fig. 4.2 is complete.

4.5.3 CBM Performance Matrix and Weighted Aggregation

The complementary side of the architecture, grouped in Fig. 4.2 under Performance Assessment, is the performance matrix to which these weights are applied, already assembled in Section 4.4 through the numerical operationalisation of Section 3.3.4. Tab. 4.13 reports the twelve decision variables whose ratings do not depend on CRM market conditions, together with the baseline ratings of the six that do; Tab. 4.14 extends the latter six to the Low and High deviations. Together, the two tables constitute the CBM performance matrix P in full and are the second input to the aggregation procedure of Section 4.5.1, requiring no further construction at this stage.

Under the baseline configuration, the weighted score attributed to each archetype is the sum, across all eighteen decision variables, of the product between each variable's global weight and its corresponding entry in Tab. 4.13 (Equations 3.17-3.18). The weights themselves were elicited from the experts as described in Section 3.5 and are reported and applied in Chapter 5; this chapter establishes only that every quantity the two equations require, the global weight of each decision variable and its harmonised, normalised performance score for each archetype, is available in a form specific to the eighteen variables and three archetypes defined for this thesis.

The same applies to the Low and High CRM market conditions, using Tab. 4.14 in place of Tab. 4.13.

4.5.4 Scenario Integration and DSS Output

The global weights W_i and the baseline performance matrix P are the inputs to the scenario module: the mechanism through which the same weighted-scoring logic is re-applied under the Low and High CRM market conditions defined in Section 3.4.1.

Section 4.4.3 already established which six of the eighteen decision variables are scenario-sensitive, as formalised by the indicator c_i in Equation 3.20, and rated their re-scored values under the Low and High deviations in Tab. 4.14. These six variables cut across four of the six macro-categories rather than belonging to a single one, since scenario sensitivity is defined in Section 3.4.2 by a variable's causal dependence on the CRM market rather than by its categorical assignment. The remaining twelve variables, without a documented mechanism linking them to that market, retain their baseline rating under both deviations. The scenario-specific performance matrix $P_{i,j,s}$ required by Equation 3.19 is therefore available in full for this thesis, without further construction.

Applying Equation 3.19 separately to the Low and High deviations, followed by Equation 3.21, identifies a scenario-specific recommended archetype in each case, mirroring the use of Equations 3.17 and 3.18 to identify the baseline recommendation in Section 4.5.3. The recommended archetypes are computed in Chapter 5, where the global weights W_i derived from the expert elicitation of Section 3.5 are applied; no separate set of weights is required for the Low and High deviations, since the weights entering Equation 3.19 are, by construction, the same W_i held fixed across all market conditions.

As set out in Section 3.4.2, the three resulting recommendations, one for the baseline configuration and one for each deviation, either coincide, making the recommendation robust to this exogenous condition, or diverge, in which case the pattern of divergence, including which archetype gains ground as the market eases or tightens, is itself informative (Schoemaker, 1995). The choice among the three archetypes, whose choice governs, among other things, whether and how the recovery of embedded CRMs remains economically and organisationally viable, may either confirm the baseline recommendation or expose its sensitivity to the CRM market condition under which the choice is made.

4.5.5 Software Implementation

The architecture described in Sections 4.5.1 to 4.5.4 is implemented as a Python pipeline that reads a single Excel workbook. The workbook contains one sheet for each expert, in which the upper triangle of the seven pairwise comparison matrices is entered, and a sheet with the normalised performance scores of Tabs. 4.13 and 4.14 for the baseline and the two scenarios. The pipeline executes the steps of Fig. 4.2 in sequence: it derives the priority vector and the consistency ratio of each matrix (Equations 3.7 to 3.13), aggregates the judgements of the experts through the geometric mean (Equation 3.9), checks the consistency of the group matrices, computes local and global weights (Equation 3.10), and applies the weighted scoring rule to the baseline and to each scenario (Equations 3.17 to 3.21). Its outputs are the consistency ratios, the weights, and the scores and rankings for each market condition, together with the category-level profiles used in the figures of this thesis. Nothing is stored between runs, so any change in a judgement or in a rating produces a complete recalculation, and the number of experts is not fixed in the code. The analysis can

therefore be repeated with different experts, weights or ratings without modifying the software.

5. Case Study: Offshore Wind Sector and NdFeB Magnet Recovery

5.1 Case Definition and Decision Context

This chapter applies the DSS developed in Chapters 3 and 4 to an illustrative case, i.e. the offshore wind case. This case does not represent a particular company or turbine model and does not aim at a sector-wide empirical conclusion. Its purpose is to demonstrate how the framework operates in a CRM-intensive industrial context.

Offshore wind is a particularly suitable context for examining the interaction between CBM choice and CRM exposure because permanent-magnet configurations are substantially more prevalent offshore than onshore. In 2022, permanent-magnet turbines accounted for 100% of offshore installations in the EU and 92% worldwide, compared with 21% and 25%, respectively, of onshore installations (Carrara et al., 2025; Ierides et al., 2025). These shares refer to the installations of a single year and do not imply that other generator designs are absent from offshore wind.

The CRM dependence of these generators defines the focus of the case. Significant dependencies on CRMs, including neodymium, dysprosium, praseodymium and terbium, are identified in the supply chain of high-performance wind-turbine magnets, together with vulnerabilities along the associated supply chains, and recycling, circularity and substitution are indicated among the relevant responses (Carrara et al., 2025). Alternative generator technologies and material-substitution options are also available, but they involve different trade-offs in terms of performance, lifetime, cost and technological maturity compared with PMSG-based configurations (Ierides et al., 2025).

The system boundary is set at the level of the offshore wind turbine, which is the common asset to which the three CBM archetypes are applied and against which their ownership, lifecycle responsibility, maintenance and end-of-life implications are compared. Within this boundary, the PMSG is the CRM-critical analytical subsystem and the NdFeB magnets constitute the specific material focus. The case does not assess the CRM footprint or material-recovery pathways of the turbine as a whole; other turbine components and material streams therefore remain outside the analytical focus. The three CBM archetypes are thus compared on the same asset and from the same OEM perspective, while the implications for NdFeB exposure and recovery are examined through the PMSG.

The focal decision-maker is the turbine OEM and the customer in the case is the wind-farm owner. Ownership and operation of the wider wind farm are not assumed to be vested in the same actor; the case focuses on the ownership and lifecycle responsibilities specified for the turbine under each CBM archetype.

The three CBM archetypes are instantiated at the level of the offshore wind turbine, using the reference configurations established in Section 4.1. The purpose is not to redefine them, but to specify how their structural differences appear in this case and how they allocate control over the PMSG and its NdFeB magnets across the asset lifecycle.

Under the product-oriented CBM, the OEM sells the turbine and ownership passes to the wind-farm owner at the point of sale, so that the PMSG and its magnets become part of a customer-owned asset. Maintenance, warranty, service and take-back arrangements can be contracted separately (Tukker, 2004; Reim et al., 2015). Since ownership is transferred, control over the PMSG at EOL rests with the wind-farm owner, and the magnets reach the OEM only if the owner uses the available take-back mechanism.

Under the use-oriented CBM, the OEM retains ownership of the turbine and grants the wind-farm owner the right to use it over a contract term in exchange for recurring payments. The wind-farm owner retains responsibility for the wider wind farm, while maintenance and lifecycle management of the turbine remain with the OEM, directly or through service providers (Tukker, 2004; Reim et al., 2015). The PMSG returns to the OEM's physical control with the return of the turbine at contract end.

Under the result-oriented CBM, the OEM retains ownership of the turbine and assumes contractual responsibility for an agreed functional or performance result delivered through it (Tukker, 2004; Baines et al., 2007; Reim et al., 2015). In the case, the OEM operates and maintains the turbines and decides on the technical means and on the lifecycle of the asset, whereas the wind-farm owner retains responsibility for the wider wind farm within which the result is delivered, without owning or operating the turbines. The PMSG therefore remains under OEM control throughout the contractual period, as ownership of the turbine is retained.

Contractual performance responsibility at turbine level is documented in the wind-energy literature, where performance-based contracts have been analysed in relation to turbine reliability and maintenance (Jin et al., 2012), with system availability and reliability incorporated into the maintenance decision (Wang et al., 2019) and, for offshore turbines specifically, with power generation as the performance standard for incentives and penalties between manufacturers and wind-farm operators (Tang et al., 2026). These studies support the formulation of contractual performance responsibility at the wind-turbine level, but do not establish that the complete retained-ownership result-oriented configuration used here is a standard commercial model in offshore wind. Moreover, the outcome-based contracts identified in the literature are predominantly availability-based, with limited clear evidence of economic-result contracts in practice (Grubic & Jennions, 2018). The result-oriented CBM is therefore treated here as a stylised reference configuration, consistent with the archetype defined in Section 4.1.

5.2 Expert Elicitation Outcomes and Criteria Weights

This section reports the outcomes of the expert elicitation described in Section 3.5 and completes the DSS with the criteria weights that the case requires. All weights, consistency ratios and scores reported in this chapter were computed with the DSS's purpose-built implementation, a Python pipeline operating on a structured Excel workbook (Section 4.5.5), which operationalises the AHP and scenario-scoring procedure of Sections 3.3 and 3.4. The implementation takes the pairwise judgements recorded during the interviews (Appendix B) as input and produces the priority vectors, consistency ratios, and scenario scores reported in this chapter directly from that input, so that every figure in this chapter is generated from the same underlying judgements. Section 5.2.1 assesses whether the pairwise judgements collected from the four experts are sufficiently coherent to be used, while Section 5.2.2

presents the resulting weights, the degree of agreement among the experts and the reading of the weighting pattern. The weights derived here enter the scoring of the CBM archetypes in Sections 5.3 and 5.4 without further modification.

5.2.1 Consistency of Expert Judgements

Each of the four experts (E1-E4, as labelled in Tab. 3.3) completed seven comparison matrices, one for the six macro-categories and one for each of the six categories, for a total of 28 matrices. The four matrices of Organisational/Capabilities are excluded from the consistency assessment, since two elements admit a single comparison and the matrix is consistent by construction (RI = 0.00 in Tab. 3.2); they are marked n/a in Tab. 5.1. The assessment therefore concerns 24 matrices, whose CR, rounded to three decimals, are reported in Tab. 5.1 together with those of the group matrices. Where a matrix was revised, the table shows the CR at first computation followed by the CR after revision. The matrix-level detail for each expert is reported in Appendix B.

Matrix	E1	E2	E3	E4	Group
Level 1 (six categories)	0.157 → 0.088	0.513 → 0.072	0.066	0.068	0.023
Market	0.000	0.016	0.069	0.033	0.006
Economic/ Business Case	0.249 → 0.040	0.022	0.243 → 0.001	0.037	0.008
Technical/ Product	0.652 → 0.011	0.056	0.117 → 0.000	0.056	0.001
Organisational/ Capabilities	n/a	n/a	n/a	n/a	n/a
Value Chain/ Ecosystem	0.011	0.025	0.033	0.006	0.003
Regulatory/ Institutional	0.270 → 0.008	0.033	0.056	0.033	0.007

Tab. 5.1 - CRs of the individual and group comparison matrices.

At first computation, seven of the 24 matrices exceeded the threshold of 0.10: the Level 1, Economic/Business Case, Technical/Product and Regulatory/Institutional matrices of E1, the Level 1 matrix of E2, and the Economic/Business Case and Technical/Product matrices of E3. E4 produced no matrix above the threshold. The excess ranged from CR = 0.117 (E3, Technical/Product) to CR = 0.652 (E1, Technical/Product). All seven matrices were reviewed with the respective expert following the procedure of Section 3.5.1 and once revised, none remained above the threshold. In six cases the review involved a single pairwise judgement, whereas E2 reconsidered the Level 1 comparison. The highest CR after revision is 0.088 (E1, Level 1), and the highest among the 17 matrices that required no revision is 0.069 (E3, Market).

Revisions are spread across the hierarchy rather than confined to one part of it. They concern the Level 1 matrix (two experts), the Economic/Business Case matrix (two), the Technical/Product matrix (two) and the Regulatory/Institutional matrix (one), whereas the Market and Value Chain/Ecosystem matrices required none. Four of the seven revisions belong to E1.

The group matrices obtained through the aggregation of Section 3.3.2 are those from which the weights are derived. Their CR is at most 0.023 (Level 1) and below 0.010 for every category-level matrix, so the second consistency check of Section 3.3.3 is also satisfied. Consistency should not be read as agreement: a matrix can be perfectly coherent and still express priorities that differ markedly from those of another expert. The extent to which the four experts converge on the same weights is examined in Section 5.2.2.

5.2.2 Weights, Expert Agreement and Interpretation

Tab. 5.2 reports the category weights W_c obtained from the group Level 1 matrix, alongside the weights that each expert's own matrix would yield. The individual weights are shown only to make agreement visible, since the DSS relies exclusively on the group weights. Because all four experts were instructed to assess criterion importance for the same reference decision problem, from the OEM perspective rather than from their own institutional positions, their individual judgements were treated as alternative assessments of a common preference structure. These judgements were therefore aggregated through the geometric mean to derive the group weights used by the DSS. The group weight of every category lies within the range spanned by the individual weights.

Category	Group (%)	E1 (%)	E2 (%)	E3 (%)	E4 (%)	Min-max (%)
Economic/Business Case	48.4	41.5	46.1	50.0	41.3	41.3-50.0
Regulatory/Institutional	15.3	18.7	4.9	23.9	15.8	4.9-23.9
Market	13.7	3.7	23.0	9.4	28.2	3.7-28.2
Technical/Product	8.5	15.5	10.2	4.8	4.2	4.2-15.5
Organisational/Capabilities	8.0	3.8	9.1	8.9	8.1	3.8-9.1
Value Chain/Ecosystem	6.1	16.8	6.6	3.1	2.4	2.4-16.8

Tab. 5.2 - Category weights: group weights and weights derived from each expert's Level 1 matrix.

Economic/Business Case is the dominant category, with a group weight of 48.4%, and it is the only one on which the experts agree both in rank and in magnitude: it ranks first for all four experts, whose individual weights range from 41.3% to 50.0%. Regulatory/Institutional and Market follow, while Technical/Product, Organisational/Capabilities and Value Chain/Ecosystem lie within 2.5 percentage points of one another (Fig. 5.1). All four experts share the view that economic considerations come first, as two of them made explicit in the interviews. E3 stated that “economic always comes out on top in all cases, it's just a question of how much”, and E1 traced the priority to the way decisions are ultimately taken in the

sector: “Unfortunately, it all comes into money. And if it’s not [...] financially viable, then it’s not getting done”.




Fig. 5.1 - Group weights of the six macro-categories.

Regulatory/Institutional, the second-ranked category, reflects a compliance-driven logic that E3 summarised as “if regulations say they have to do it, they’ll do it. If it doesn’t say they have to do it, they won’t” (E3). Agreement on the remaining categories is considerably weaker. Market ranges from 3.7% (E1, ranked last) to 28.2% (E4, ranked second), Regulatory/Institutional from 4.9% (E2, ranked last) to 23.9% (E3, ranked second), and Value Chain/Ecosystem from 2.4% (E4) to 16.8% (E1). The weight of Technical/Product is highest for E1 (15.5%) and below 5% for E3 and E4. The weights of the categories other than Economic/Business Case should therefore be regarded as the outcome of a panel with heterogeneous priorities, and their influence on the recommendation is tested through the individual weight profiles in Section 5.5.

Tab. 5.3 reports the local and global weights of the 18 decision variables, ordered by global weight. The global weight of each variable is the product of its category weight and its local weight (Equation 3.10) and is the value applied in the scoring of Sections 5.3 and 5.4.

Decision variable	Macro-category	Local weight (%)	Global weight (%)
Operational Service-Delivery Cost Burden	Economic/Business Case	33.4	16.1
Provider Investment/Capital Requirement	Economic/Business Case	32.9	15.9
Value Creation and Capture Potential	Economic/Business Case	18.8	9.1

Decision variable	Macro-category	Local weight (%)	Global weight (%)
Customer Acceptance of the CBM's Ownership/Access Arrangement	Market	60.4	8.3
CRM Price Volatility Exposure	Economic/Business Case	14.9	7.2
Regulatory and Financial Incentive Availability	Regulatory/Institutional	44.3	6.8
Exposure to Regulatory Change over the Contract Horizon	Regulatory/Institutional	36.7	5.6
Service and Maintenance Capability	Organisational/Capabilities	66.3	5.3
Product Reliability and Durability	Technical/Product	61.5	5.3
Lifecycle Data Accessibility and Control	Value Chain/Ecosystem	64.7	3.9
Customer Asset-Care and Moral-Hazard Exposure	Market	23.4	3.2
Regulatory Traceability Compliance Capability	Regulatory/Institutional	19.0	2.9
Organisational Restructuring Requirement	Organisational/Capabilities	33.7	2.7
Demand/Usage Forecasting Dependency	Market	16.3	2.2
Technology and Material Substitution Risk over the Contract Horizon	Technical/Product	19.8	1.7
Technological Infrastructure and Complexity Requirement	Technical/Product	18.7	1.6
Ecosystem Coordination Complexity	Value Chain/Ecosystem	18.4	1.1
CRM Recovery and Asset Return Effectiveness	Value Chain/Ecosystem	16.8	1.0

Tab. 5.3 - Local and global weights of the 18 decision variables.

The global weight is heavily concentrated. The three highest-ranked variables, all belonging to Economic/Business Case, account for 41.1% of the total and the highest five for 56.6%, whereas the five lowest-ranked variables together account for less than 8%. This concentration anticipates the fully compensatory nature of the weighted-sum aggregation discussed in Chapter 6 (Kaliszewski & Podkopaev, 2016): under an additive rule, a small number of high-weight variables can determine the ranking regardless of how the

alternatives perform on the remaining criteria. Operational Service-Delivery Cost Burden (16.1%) and Provider Investment/Capital Requirement (15.9%) are practically tied, and their difference of 0.2 percentage points is too small for their order to be meaningful. Operational Service-Delivery Cost Burden is nevertheless the only variable ranked among the top three by each of the four experts. Within the category itself, the experts do not agree on a leading variable: Provider Investment/Capital Requirement leads for E1 (jointly with CRM Price Volatility Exposure) and for E4, Value Creation and Capture Potential for E2, and Operational Service-Delivery Cost Burden for E3.

Local weights show a different degree of consensus. Customer Acceptance of the CBM's Ownership/Access Arrangement is the leading variable of Market for all four experts (group local weight 60.4%), and Lifecycle Data Accessibility and Control is the leading variable of Value Chain/Ecosystem for all four (64.7%). Service and Maintenance Capability leads Organisational/Capabilities for three experts, E3 being the exception. The leadership of the other two categories is split: in Technical/Product, Product Reliability and Durability leads at group level (61.5%) and for two experts, while in Regulatory/Institutional Regulatory and Financial Incentive Availability leads at group level (44.3%), ahead of Exposure to Regulatory Change over the Contract Horizon (36.7%), with two experts on each side.

The prominence of Lifecycle Data Accessibility and Control within Value Chain/Ecosystem is consistent with the information barriers reported in the interviews. E2 recalled that “the OEMs were really reluctant to share any warranty information”, while E3 observed that “operators can buy repaired, refurbished, remanufactured components or systems, but the procurement departments don't necessarily know what the difference is between them [...] it's very difficult for the supplier to be clear what the differences are”. Both remarks point to limited data sharing and to unclear distinctions between recovery routes, which may explain why the experts agree on the importance of data control within the category while assigning limited weight to the category.

Because the category weight scales every local weight, the ranking of the variables is driven largely by the category to which they belong. CRM Price Volatility Exposure has the lowest local weight in Economic/Business Case (14.9%), yet its global weight (7.2%) exceeds that of every variable outside this category except Customer Acceptance of the CBM's Ownership/Access Arrangement. Conversely, Lifecycle Data Accessibility and Control holds the second-highest local weight among the 18 variables (64.7%, after Service and Maintenance Capability with 66.3%) but ranks only tenth globally. At the bottom of the ranking, CRM Recovery and Asset Return Effectiveness has the lowest global weight (1.0%, from a local weight of 16.8%). This outcome reflects the judgements of E2, E3 and E4, who rank the variable between sixteenth and eighteenth; only E1 places it considerably higher, in eighth position. It should be read as the importance of this criterion for the OEM's choice among CBM archetypes, not as the importance of recovery as an objective: the recovery mechanism also enters the model through CRM Price Volatility Exposure and Value Creation and Capture Potential (Section 4.3), which together carry 16.3% of the total weight.

5.3 Baseline CBM Ranking

This section applies Equations 3.17 and 3.18 to the global weights of Tab. 5.3 and to the baseline CBM performance matrix (Tab. 4.13) and reports the resulting ranking of the three CBM archetypes. It presents the overall scores and the recommended archetype, the

performance of each archetype within the macro-categories, and the contribution of the macro-categories to the overall scores and the decision variables that determine the ranking.

Under the baseline condition, the product-oriented CBM obtains the highest overall score (0.781), followed by the use-oriented CBM (0.704) and the result-oriented CBM (0.589), as shown in Fig. 5.2. The DSS therefore recommends the product-oriented CBM (Equation 3.18). The gap between the first and the second archetype is 0.077, and that between the second and the third is 0.115.




Fig. 5.2 - Overall scores of the three CBM archetypes under the baseline condition.

Fig. 5.3 shows the performance of each CBM archetype within each macro-category, on a scale from 0 to 1 in which higher values are more favourable. It is built from the same normalised scores as Fig. 4.1 (Section 4.4.4), but the two figures play different roles. Fig. 4.1 reported the simple average of the scores of each category, with no importance weights, and was descriptive: it was not part of the DSS scoring. Fig. 5.3 is the first step of that scoring, since it combines the scores through the local weights of Tab. 5.3, so that within each category the decision variables that the panel considers more important count more. It therefore shows how each archetype performs in each category as seen through the priorities of the panel within that category. The weights of the macro-categories are not applied in this figure, and it does not reflect how much each category weighs in the overall score; this is introduced in Fig. 5.4.




Fig. 5.3 - Performance of the three CBM archetypes by macro-category under the baseline condition.

No archetype is ahead in every category. The product-oriented CBM reaches the maximum value in Market and in Organisational/Capabilities (1.00 in both) and the highest value in Economic/Business Case (0.84) but is lowest in Value Chain/Ecosystem (0.33), Regulatory/Institutional (0.58) and Technical/Product (0.59). The result-oriented CBM shows the opposite profile, with the highest values in Value Chain/Ecosystem (0.88) and Regulatory/Institutional (0.75) and the lowest in Market (0.41), Economic/Business Case (0.56) and Organisational/Capabilities (0.33). The use-oriented CBM occupies an intermediate position in five of the six categories and shares the highest value in Technical/Product with the result-oriented CBM (0.81 for both). The local weights alter two profiles with respect to Fig. 4.1. In Technical/Product, the product-oriented CBM moves from the highest unweighted average (0.78) to the lowest value (0.59), because most of the local weight falls on Product Reliability and Durability (61.5%), where it is rated least favourably. In Economic/Business Case, the unweighted averages of the product- and use-oriented CBMs coincided (0.75), whereas the weighted values separate them (0.84 and 0.73), since Provider Investment/Capital Requirement and Operational Service-Delivery Cost Burden hold 66.3% of the local weight of the category. The relatively narrow spread of the overall scores reflects two features of the scoring structure rather than an even three-way trade-off. First, the normalised scale is bounded between 0.33 and 1.00 (Section 3.3.4), which caps how far any two archetypes' scores can diverge. Second, because most decision variables move monotonically along the same ownership-retention continuum that orders the three archetypes themselves (Product, Use, Result; Section 3.2.1), the use-oriented CBM occupies an intermediate position on nearly all of them, and its overall score falls, almost mechanically, between the other two. Despite this narrow spread, the product-oriented CBM still leads outright in three of the six macro-categories (Market, Economic/Business Case,

Organisational/Capabilities) against two led by the result-oriented CBM (Value Chain/Ecosystem, Regulatory/Institutional; Fig. 5.3), but its overall lead reflects the greater combined weight of its strongest categories (70.1%, against 21.4% for the result-oriented CBM's) rather than a wide margin in any single one of them.

Fig. 5.4 combines the profiles of Fig. 5.3 with the weights of the macro-categories (Tab. 5.2). Each cell reports the contribution of a macro-category to the overall score of an archetype, that is, the weight of the category multiplied by the score plotted in Fig. 5.3, so that each column adds up to the corresponding overall score of Fig. 5.2 (up to rounding).




Fig. 5.4 - Contribution of each macro-category to the overall score of the three CBM archetypes under the baseline condition.

Economic/Business Case is the dominant row: it supplies 0.405 of the overall score of the product-oriented CBM, 0.354 of that of the use-oriented CBM and 0.269 of that of the result-oriented CBM. Together with Market, it is also the row in which the contributions of the three archetypes are most distant from one another. In the remaining four rows, which carry 37.9% of the total weight, the contributions are lower in absolute terms and, with the exception of Organisational/Capabilities, are higher for the use- and result-oriented CBMs. Since the macro-categories enter the score with very different weights, the strengths of these two archetypes in Value Chain/Ecosystem and Regulatory/Institutional translate into contributions that remain modest next to that of Economic/Business Case.

Within Economic/Business Case, the ranking is decided by two decision variables. Provider Investment/Capital Requirement and Operational Service-Delivery Cost Burden jointly carry 32.0% of the total weight and are rated most favourably for the product-oriented CBM (1.00, against 0.67 for the use-oriented and 0.33 for the result-oriented CBM). Together, they contribute 0.106 to the lead of the product-oriented CBM over the use-oriented CBM and 0.215 to its lead over the result-oriented CBM, the latter exceeding the total gap of 0.192. On the remaining sixteen variables, both the use-oriented and the result-oriented CBM are ahead of the product-oriented CBM. The product-oriented CBM therefore prevails because

of the two variables on which the panel places the greatest weight, and not because of a broadly more favourable performance across the framework. The same two variables account for 0.109 of the 0.115 that separates the use-oriented from the result-oriented CBM, which explains why these two follow the same order. This outcome is consistent with the literature on CBM adoption, which identifies financial conditions among the main obstacles faced by firms, from the large upfront investments that CBMs require to the uncertainty of their returns (Hina et al., 2022; Geissdoerfer et al., 2023).

This has a direct implication for the focus of the case. The product-oriented CBM is rated lowest (0.33) on CRM Recovery and Asset Return Effectiveness, Lifecycle Data Accessibility and Control, Regulatory Traceability Compliance Capability and CRM Price Volatility Exposure, the four decision variables most directly tied to the retention of the asset and of the NdFeB magnets it contains. On these variables, which carry 15.1% of the total weight, the result-oriented CBM is rated 1.00 and would lead by 0.101. This advantage reflects the link between ownership retention and control over the asset at the end of its life, which keeps the materials within the boundaries of the OEM (Bocken et al., 2016; Cimprich et al., 2023). Its counterpart is an operating burden, since return logistics and labour-based remanufacturing can cost more than the material savings they generate (Hina et al., 2022). The baseline recommendation is therefore not a finding that the product-oriented CBM is the most effective route to NdFeB recovery. It indicates that, for the panel, the advantage of retaining the asset in terms of return, control and material exposure is outweighed by the capital commitment and operating burden that retention entails for the OEM.

5.4 CBM Ranking under Alternative CRM Market Conditions

The scenario mechanism defined in Section 3.4 is carried into the case by interpreting the generic Low and High CRM market scenarios specifically for the CRMs relevant to the NdFeB magnets used in the selected PMSG configuration. The baseline remains the neutral intermediate condition established in Section 4.4. The Low scenario therefore represents a relatively lower-price and more predictable market condition, whereas the High scenario represents a relatively higher-price and less predictable market condition.

The price dimension of this distinction has a direct historical basis. The prices of neodymium, praseodymium, dysprosium and terbium varied substantially between 2005 and 2015, with a pronounced increase in 2011 following Chinese export restrictions and a subsequent decline after 2013 (Pavel et al., 2017). More directly, the recurrence of the 2011 price spike has been modelled as a historical price-event scenario for NdFeB permanent magnets in direct-drive wind-turbine generators, showing a substantial effect on generator cost (Leader et al., 2019). This evidence provides an empirical basis for distinguishing lower- and higher-price states relevant to the case. The predictability dimension remains a qualitative scenario assumption, as set out in Section 3.4, rather than a numerical estimate of future NdFeB prices. The two scenarios are therefore used as contrasting illustrative market states against which the same DSS is re-applied, rather than as forecasts of future material prices.

In both scenarios, the global weights of Tab. 5.3 are kept fixed and only the scores of the six scenario-sensitive decision variables are updated (Tab. 4.14). Section 5.4.1 reports the ranking under the Low scenario and Section 5.4.2 under the High scenario, using for each the layout of Figs. 5.2 to 5.4.

5.4.1 Low CRM Market Scenario

Under the Low scenario, four of the six scenario-sensitive decision variables change with respect to the baseline. Operational Service-Delivery Cost Burden and CRM Price Volatility Exposure become more favourable, Technology and Material Substitution Risk over the Contract Horizon eases for the use- and result-oriented CBMs, and Value Creation and Capture Potential declines for the same two archetypes, since the recoverable material is worth less. Customer Acceptance and Service and Maintenance Capability keep their baseline scores.

Fig. 5.5 reports the resulting overall scores. The ranking is unchanged: the product-oriented CBM remains first (0.805), followed by the use-oriented CBM (0.757) and the result-oriented CBM (0.620), and all three scores are higher than in the baseline. The increase is uneven, since it amounts to 0.024 for the product-oriented CBM, 0.053 for the use-oriented CBM and 0.031 for the result-oriented CBM. The lead of the product-oriented CBM over the use-oriented CBM therefore narrows from 0.077 to 0.048, whereas its lead over the result-oriented CBM remains almost unchanged (from 0.192 to 0.185).




Fig. 5.5 - Overall scores of the three CBM archetypes under the Low CRM market scenario.

Two effects account for most of the narrowing of the gap with the use-oriented CBM. The operating cost burden of the use-oriented CBM eases to the level of the product-oriented CBM, which removes an advantage of 0.053 that the latter held on Operational Service-Delivery Cost Burden. At the same time, the value that the use-oriented CBM derives from recovered material declines to the level of the product-oriented CBM, which offsets 0.030 of that reduction. CRM Price Volatility Exposure improves by one step for both archetypes and leaves the gap almost unchanged. The advantage of the product-oriented CBM on Provider Investment/Capital Requirement, which is not scenario-sensitive, is not affected and, together with Market and Organisational/Capabilities, keeps it first.




Fig. 5.6 - Performance of the three CBM archetypes by macro-category under the Low CRM market scenario.

Since the scenario acts only on variables belonging to Economic/Business Case and Technical/Product, the profiles of Fig. 5.6 differ from the baseline in these two categories alone. In Economic/Business Case, the values rise to 0.89 for the product-oriented CBM, 0.83 for the use-oriented CBM and 0.61 for the result-oriented CBM, from 0.84, 0.73 and 0.56 in the baseline. In Technical/Product, the use- and result-oriented CBMs move from 0.81 to 0.87, while the product-oriented CBM remains at 0.59. The contributions of Fig. 5.7 show the same pattern: the difference between the product- and use-oriented CBMs in Economic/Business Case falls from 0.051 to 0.029 (0.430 and 0.401), and all other rows are identical to those of Fig. 5.4.




Fig. 5.7 - Contribution of each macro-category to the overall score of the three CBM archetypes under the Low CRM market scenario.

The Low scenario therefore confirms the recommendation of the product-oriented CBM, with a smaller margin over the use-oriented CBM. In the case, the decline of Value Creation and Capture Potential reflects the lower market value of the NdFeB material recoverable at the end of the life of the turbines, which reduces the advantage that ownership retention gives to the use- and result-oriented CBMs, while the easing of price pressure lowers the operating burden that these two archetypes carry.

5.4.2 High CRM Market Scenario

Under the High scenario, five of the six scenario-sensitive decision variables change with respect to the baseline. Operational Service-Delivery Cost Burden and Service and Maintenance Capability become less favourable for the product- and use-oriented CBMs, Value Creation and Capture Potential declines for the product-oriented CBM, Technology and Material Substitution Risk over the Contract Horizon becomes less favourable for the product- and result-oriented CBMs, and Customer Acceptance falls for the product-oriented CBM while rising for the use- and result-oriented CBMs. CRM Price Volatility Exposure keeps its baseline scores, since the ratings of the use- and result-oriented CBMs are held by assumption (Section 4.4.3).

Fig. 5.8 shows that the ranking changes. The use-oriented CBM becomes first (0.658), ahead of the product-oriented CBM (0.646) and the result-oriented CBM (0.611). With respect to the baseline, the score of the product-oriented CBM falls by 0.135 and that of the use-oriented CBM by 0.046, whereas the score of the result-oriented CBM rises by 0.022. The lead of the use-oriented CBM over the product-oriented CBM is 0.012, and the distance between the product- and result-oriented CBMs, which was 0.192 in the baseline, becomes 0.035. The margin between the first two archetypes is small, and the reversal should be read as an indication that their order changes under this condition, not as a clear-cut superiority of the use-oriented CBM.




Fig. 5.8 - Overall scores of the three CBM archetypes under the High CRM market scenario.

The reversal is produced mainly by two variables on which the two archetypes move in different directions, and which shift the gap between them from +0.077 to -0.012. Customer Acceptance accounts for 0.055 of this change, since it falls for the product-oriented CBM and rises for the use-oriented CBM by the same step. Value Creation and Capture Potential, which declines only for the product-oriented CBM, accounts for 0.031. By contrast, Operational Service-Delivery Cost Burden and Service and Maintenance Capability worsen by one step for both archetypes, and their effect on the gap is close to zero, although they account for a large part of the loss of both scores. The result-oriented CBM gains 0.028 from Customer Acceptance and loses 0.006 on substitution risk, which explains its net increase of 0.022.




Fig. 5.9 - Performance of the three CBM archetypes by macro-category under the High CRM market scenario.

The profiles of Fig. 5.9 change in four categories. Market is the category in which the archetypes converge: the product-oriented CBM falls from 1.00 to 0.80 and the use-oriented CBM rises from 0.59 to 0.79, so that their contributions in Fig. 5.10 become almost equal (0.110 and 0.108, against 0.137 and 0.081 in the baseline). In Economic/Business Case, the values of the product- and use-oriented CBMs decline to 0.66 and 0.62 while that of the result-oriented CBM stays at 0.56, and the contributions narrow to 0.321, 0.299 and 0.269. In Organisational/Capabilities, the product-oriented CBM falls from 1.00 to 0.78 and the use-oriented CBM from 0.67 to 0.44, and in Technical/Product the product- and result-oriented CBMs decline to 0.52 and 0.74. Value Chain/Ecosystem and Regulatory/Institutional, which contain no scenario-sensitive variable, are identical to the baseline.




Fig. 5.10 - Contribution of each macro-category to the overall score of the three CBM archetypes under the High CRM market scenario.

The High scenario acts mainly on the customer side and on the cost of service. As prices rise and supply becomes less predictable, customers become more willing to accept access-based arrangements, whereas the product-oriented CBM, which does not retain the recovered NdFeB material, loses the value of that material and sees its warranty and service commitments become more onerous. The result-oriented CBM benefits from the change in customer acceptance but remains last, since the two variables on which it is weakest, Provider Investment/Capital Requirement and Operational Service-Delivery Cost Burden, do not improve: the former is not scenario-sensitive and the latter is already at its lowest score in the baseline.

5.5 Sensitivity to the Weights of the Individual Experts

Section 5.2.2 showed that the experts agree on the dominance of Economic/Business Case but diverge on the other categories, and that the DSS relies on the group weights. This section tests whether the baseline recommendation depends on the aggregation of the four judgements. Each expert's own weight profile, obtained from the matrices of that expert alone with the same procedure used for the group weights (Section 3.3.2), replaces the group weights in equation (3.17), while the baseline performance matrix of Tab. 4.13 is kept unchanged. No new data are needed. The resulting scores are reported in Tab. 5.4.

Weight profile	Product-oriented	Use-oriented	Result-oriented	Ranking
Group (Baseline)	0.781	0.704	0.589	P > U > R
E1	0.672	0.673	0.689	R > U > P
E2	0.764	0.726	0.668	P > U > R
E3	0.784	0.752	0.603	P > U > R
E4	0.833	0.686	0.527	P > U > R

Tab. 5.4 - Overall scores of the three CBM archetypes obtained with the weight profile of each expert (Note: P, U and R denote the product-, use- and result-oriented CBM).

The product-oriented CBM ranks first with the profiles of E2, E3 and E4, with leads over the use-oriented CBM of 0.038, 0.032 and 0.147. With the profile of E1 the ranking changes: the result-oriented CBM is first (0.689), and the product- and use-oriented CBMs are practically tied (0.672 and 0.673, a difference below 0.001). The difference is due to two features of the profile of E1. The first is the weight of Value Chain/Ecosystem (16.8%, against 6.1% for the group), the category in which the product-oriented CBM is rated lowest and the result-oriented CBM highest. The second is the lower weight given to Provider Investment/Capital Requirement and Operational Service-Delivery Cost Burden, which together carry 24.5% in the profile of E1 and 32.0% in the group, and which are the variables that sustain the lead of the product-oriented CBM in the baseline.

The recommendation of the product-oriented CBM is therefore supported by three of the four expert profiles, but it is not unanimous, and the aggregation into group weights does hide a divergence that reaches the first place. This confirms that the weights of the categories other than Economic/Business Case, on which the panel is most heterogeneous, matter for the outcome. Read together with Section 5.4, the two checks indicate that the lead of the product-oriented CBM is real but not unconditional: it depends both on a favourable rating on capital requirement and operating cost and on a weighting of the panel in which Value Chain/Ecosystem and Technical/Product do not carry more than a small share. This dependence is consistent with the role assigned to criterion weighting in the framework: since the DSS is designed to reflect the priorities of the OEM for which it is applied (Section 3.2.1), a recommendation that shifts with the weighting profile is the expected behaviour of a preference-sensitive instrument rather than evidence against the reliability of the underlying scoring procedure. It indicates, instead, that the choice of weighting profile is itself consequential, and should be made deliberately by the OEM applying the framework rather than treated as a fixed input.

5.6 Discussion

Sections 5.3 to 5.5 can each be read as a self-contained result: a baseline ranking, two departures from it under alternative CRM market conditions, and a check against the weighting profile of each expert. Read together, they converge on a single point rather than on three separate ones. The recommendation the DSS returns for the case is not a fixed verdict on which CBM archetype an OEM should adopt; it is the output of two conditions that the framework, by design, keeps distinct rather than folding into one number. One

condition is external to the OEM, the state of the CRM market at the time the decision is taken, examined in Section 5.4. The other is internal to it, the preference structure used to represent the OEM's priorities across the eighteen decision variables, examined in Section 5.5. What follows argues that both conditions move the baseline ranking through the same underlying tension, between the capital and cost advantage of the product-oriented CBM and the material-control advantage retained by the other two archetypes, a tension that Sections 5.3 to 5.5 each register from a different angle without naming it explicitly.

The baseline ranking of Section 5.3 sets the reference point: the product-oriented CBM leads because two variables, Provider Investment/Capital Requirement and Operational Service-Delivery Cost Burden, outweigh a broader but less heavily weighted advantage the other two archetypes hold elsewhere in the framework, in line with the capital and return barriers the CBM-adoption literature already documents (Hina et al., 2022; Assmann et al., 2023; Geissdoerfer et al., 2023). What the rest of this section asks is how far this lead travels once either of the two conditions identified above is allowed to move.

Under the Low scenario, the product-oriented CBM keeps its lead, but the margin over the use-oriented CBM narrows (Section 5.4.1). Part of this narrowing comes from the use-oriented CBM's operating-cost burden easing to match the product-oriented CBM's already most favourable rating on that variable; the rest comes from the erosion of the advantage the use- and result-oriented CBMs draw from the material they retain, which is worth less once the material itself is worth less. This second effect is the numerical counterpart of a mechanism the CBM-adoption literature usually reports from the opposite side: firms surveyed across the European Union cite the uncompetitively low price of virgin material relative to secondary material among the most critical market barriers to circular adoption (Kirchherr et al., 2018). The case reproduces the same asymmetry from the recovering archetype's side of the transaction, giving a numerical face to the more general premise that circular business models compete on value retained across the life of an asset rather than on value added in a single transaction (Lüdeke-Freund et al., 2019): the retained value the framework credits to the use- and result-oriented CBMs is only as large as the market allows it to be.

Under the High scenario, the same tension pushes the ranking past the point of reversal, by a margin too small to read as a clear preference for the use-oriented CBM over the product-oriented one (Section 5.4.2). The reversal is carried less by the operating cost that worsens for both archetypes, and barely moves the gap between them, than by a shift on the customer side: as the price and financing risk of owning the asset outright rise, customers grow more willing to accept an access-based arrangement that transfers that risk to the OEM, a pattern documented in the literature on the leasing of capital goods (Frederiksen et al., 2021; Reim et al., 2015; Tukker, 2004). Under this condition, the product-oriented CBM loses ground on both dimensions: it loses part of the customer-acceptance advantage it held at baseline, while becoming relatively less able to capture the value associated with recoverable material than the ownership-retaining archetypes. The closest precedent in the literature reviewed in Chapter 2 is GE's response to the rhenium price shock: after the shock, GE moved away from disposing of end-of-life turbine blade assemblies as general scrap and developed recovery and recycling practices aimed at retaining rhenium within its value chain (Gaustad et al., 2018; Cimprić et al., 2023). Separately, the aviation MRO model provides a precedent for the link between retained turbine ownership, lifecycle responsibility and control over

material flows (Cimprich et al., 2023). These precedents are not reproduced directly in the present case: the material, component and industry differ, while the DSS translates the underlying logic into an ex-ante comparison of CBM alternatives under a market condition that has not yet occurred. A further, more modest logic reinforces this. Under CRM price pressure, maintenance is assumed to shift from replacing worn components toward repairing and reconditioning them, a labour-for-capital trade-off documented for recovery decisions generally (Cimprich et al., 2023; Gaustad et al., 2018; Lüdeke-Freund et al., 2019). This moderates, rather than drives, the deterioration common to the product- and use-oriented CBMs under this scenario and is best read as a modelling assumption consistent with the R-hierarchy logic of Section 2.2.4, not as an independently evidenced mechanism in this specific context.

Section 5.5 moves along the second, internal axis and lands on the same tension from the other side. Under the weighting profile of expert E1, the ranking reverses, with the result-oriented CBM taking first place and the product- and use-oriented CBMs separated by less than 0.001. The reversal is driven by E1's higher weighting of Value Chain/Ecosystem, the category in which the product-oriented CBM performs worst and the result-oriented CBM best, together with a lower combined weight for Provider Investment/Capital Requirement and Operational Service-Delivery Cost Burden, the two variables that underpin the product-oriented lead in the group solution. The result shows that the ranking is sensitive to the preference structure used to represent the OEM's priorities. The reversal is therefore not a contradiction of the baseline result, but an illustration of how changes in criterion importance can alter the outcome of the same performance assessment. Since the DSS is designed to make the priorities underlying the decision explicit rather than to impose a generic weighting structure, this sensitivity is part of its intended function. The two checks are not independent robustness tests that happen to point the same way; they tilt a single balance from two different directions, even though they do so through different variables. The High scenario does not erode the product-oriented CBM's advantage on Provider Investment/Capital Requirement directly, since that variable is not scenario-sensitive, nor mainly through Operational Service-Delivery Cost Burden, whose movement affects the product- and use-oriented CBMs by a similar amount and barely shifts the gap between them; it acts instead through Customer Acceptance and the value of retained material, as set out above. The profile of E1 acts more directly on the variables that anchor the baseline lead itself, reducing the weight given to capital and operating cost while raising the weight given to Value Chain/Ecosystem. The two routes are accordingly not identical, one works through how well the product-oriented CBM performs, the other through how much its strongest variables count, but both close in on the same tension this section opened with. The baseline recommendation is therefore not fragile across many unrelated dimensions of the framework; it rests on one traceable balance, which both the market the OEM faces and the judgement of whoever applies the framework on its behalf are shown, in this case, to be capable of shifting.

The practical implication for an OEM applying the framework beyond this illustrative case follows from the fact that the ranking changes when the conditions and priorities represented in the model change. The DSS can therefore support the OEM by showing not only which archetype obtains the highest score under a given configuration, but also how the comparison changes when the CRM market condition or the importance assigned to the criteria changes.

Its value lies accordingly in making the conditions and priorities that drive the choice explicit and traceable, while leaving the final decision to managerial judgement.

6. Conclusions

This thesis addressed a specific managerial problem: how an OEM can compare Product-, Use- and Result-oriented CBMs for the recovery of CRMs under market conditions that may change over time. The offshore wind setting provides a relevant context for examining this problem. Permanent-magnet generators link turbine deployment to demand for NdFeB magnets (Carrara et al., 2025), while the materials contained in these magnets are still recovered from end-of-life products only to a limited extent (Chinwego et al., 2025; Rizos et al., 2024). Following a DSR approach, the thesis developed a scenario-based DSS and applied it to an illustrative case centred on permanent-magnet generators and NdFeB recovery in the offshore wind setting.

Existing research addresses important parts of this managerial problem but leaves the relevant perspectives only partially integrated. Criticality-assessment methods provide structured approaches to identifying material-related risks but offer limited support for linking these assessments to the subsequent selection of mitigation strategies (Otterbach & Fröhling, 2025). CBM-selection and multi-criteria decision-making approaches, in turn, provide established bases for comparing alternatives, but offer limited structured support for comparing CBM archetypes from the perspective of an OEM under changing CRM market conditions. This thesis contributes a DSS that addresses this intersection by making the comparison among CBM archetypes explicit and structured. The resulting recommendation remains conditional on both the state of the CRM market and the priorities of the OEM applying the framework, rather than being reduced to a single fixed outcome. This framework therefore not only identifies the preferred alternative under a given set of assumptions, but also makes visible which changes in external conditions or decision-maker priorities would be sufficient to alter the resulting recommendation.

Read against the three research questions set out in Chapter 1, the thesis offers the following answers. The first question, concerning the factors that should inform an OEM's choice among CBMs for CRM recovery, is answered by eighteen decision variables organised into six macro-categories: Market, Economic/Business Case, Technical/Product, Organisational/Capabilities, Value Chain/Ecosystem, and Regulatory/Institutional. These variables were derived from literature (Tukker, 2004; Reim et al., 2015; Konietzko et al., 2020a; Assmann et al., 2023).

The second question, concerning how these factors can be translated into a managerial DSS, is addressed through a two-level AHP architecture. This reduces the pairwise-comparison burden placed on each expert from 153 direct comparisons to 34 hierarchically organised comparisons (Saaty, 1980, 2008), combined with a weighted-sum scoring procedure through which criterion weights and archetype performance are aggregated into a single comparable score (Kaliszewski & Podkopaev, 2016). The architecture is operationalised as a reusable Python pipeline operating on a structured Excel workbook (Section 4.5.5).

The third question, concerning how market uncertainty affects the resulting recommendation, is addressed through the Low and High CRM market scenarios described in Chapter 5. Under the Low CRM market scenario, the product-oriented CBM remains the preferred archetype, though by a narrower margin than in the baseline; under the High

scenario this ranking reverses, narrowly favouring the use-oriented CBM. A sensitivity check using each expert's weighting profile shows that the same baseline ordering can also be reversed in favour of the result-oriented CBM through a change in decision-maker priorities. These checks indicate that the case's baseline recommendation reflects a balance between the lower capital and cost burden of the product-oriented CBM and the greater material control retained by the other two archetypes. This balance is sensitive both to external market conditions and to the OEM's decision priorities.

These findings should be read alongside some limitations inherent in the design choices underlying the framework.

The expert panel comprises four informants, none of whom is formally employed by a turbine OEM, selected purposively for their relevant professional or research experience rather than drawn as a statistically representative sample (Patton, 2002). The resulting weights should therefore be interpreted as an informed basis for demonstrating the DSS, rather than as an estimate of the preferences of a wider population of OEM decision-makers; the limited panel size constrains the generalisability of the weighting pattern reported in Chapter 5.

A second limitation concerns the ordinal granularity of the Low/Medium/High rating scale used throughout this framework. The DSS can establish the relative performance of each archetype on a given decision variable, but it cannot indicate the magnitude of that difference in real economic terms. Moreover, the three-level resolution can prevent an already extreme rating from moving further in the same direction when a scenario condition intensifies. This ceiling-and-floor effect is a limitation of the measurement scale, rather than evidence that the underlying exposure of the archetype concerned has stopped changing. A related simplification concerns the CRM market scenarios themselves: Low, Baseline and High conditions are treated, by design, as joint states of price level and predictability (Section 4.4.3), whereas an actual CRM market could in principle combine these differently, for instance remaining volatile at a comparatively low price, or stable at a comparatively high one.

A further limitation follows from a property of the AHP itself rather than from a choice specific to this application. The weighted-sum logic through which criterion weights and archetype performance are combined is, by construction, fully compensatory (Kaliszewski & Podkopaev, 2016): a markedly unfavourable rating on one decision variable can be offset by sufficiently favourable performance on others, without a minimum threshold of acceptability. This property is directly visible in the baseline case, where the product-oriented CBM receives the lowest rating among the three archetypes on the four decision variables most directly related to NdFeB recovery and retention, yet obtains the highest overall score. The aggregate score should therefore not be interpreted in isolation: the variable-level decomposition reported in Chapter 5 is necessary to identify the trade-offs underlying the recommendation. This reinforces the role of the DSS as a support for managerial judgement rather than as an automated decision rule. Introducing a minimum acceptability threshold, as in non-compensatory or partially compensatory methods such as ELECTRE or PROMETHEE, would address this at the cost of a different aggregation logic and additional threshold parameters to elicit from decision-makers. A related limitation concerns the eighteen decision variables themselves. As noted in Section 3.2.5, their consolidation addressed conceptual overlap at the design stage rather than through a statistical or correlation-based test of the independence that AHP-based scoring assumes

(Saaty, 1980, 2008); where criteria are expected to interact, the Analytic Network Process offers an established alternative within the same framework, at the cost of a substantially more demanding elicitation process. A related limitation concerns the Low/Medium/High ratings themselves (Section 3.2.6). Although grounded in the archetype-specific literature, these ratings were not separately validated with the experts, whose interviews were already focused on the AHP weighting exercise. Because the scoring procedure is compensatory, as noted above, individual ratings can still influence the final recommendation.

Finally, the case itself remains illustrative: it does not represent a specific turbine model or company. The resulting recommendation should therefore be interpreted as a demonstration of how the framework operates in a CRM-intensive offshore-wind setting, rather than as an empirical finding about the offshore wind sector.

These limitations point towards several directions for future developments. The most direct is its application within one or more OEMs, to an actual strategic decision rather than to a panel of informed proxies. This would allow both the practical usability of the DSS to be assessed and weighting patterns to be examined across OEMs with different organisational priorities. A second direction concerns applying the framework to other CRM-intensive components and sectors, such as electric-vehicle traction motors, to explore its application beyond the offshore-wind context examined in this thesis. A third direction addresses the ordinal granularity discussed above by modelling the market-predictability dimension of the CRM scenarios along a continuous rather than a three-state scale. This could involve linking the scenarios to existing geopolitical-risk indices or estimated probabilities of export restrictions. Where comparable cost and cash-flow data become available across the three archetypes, the framework could also be extended by integrating them (Gaudenzi et al., 2021). The same direction could also test a wider range of market conditions in which price level and predictability diverge. A fourth direction concerns the CBM archetypes themselves. The framework compares three distinct canonical configurations, whereas commercial practice may combine elements of more than one, such as a sale accompanied by a contractually binding take-back obligation. Extending the decision-variable set and rating logic to accommodate hybrid or intermediate configurations would therefore test whether the present structure remains adequate when the alternatives are no longer sharply delineated. A fifth direction concerns the single-decision-maker perspective adopted throughout this thesis. The case explicitly separates the OEM, as the focal decision-maker, from the customer who owns and operates the asset in which the component sits, whose priorities are not assumed to coincide with those of the OEM. Extending the framework to incorporate the customer's priorities alongside those of the OEM, whether through a parallel weighting exercise or a joint elicitation process, would move the DSS beyond a single-actor decision perspective and allow it to represent where the two parties' priorities align and where they diverge.

A final, more exploratory direction concerns the integration of artificial intelligence into the decision-support architecture. Research on AI-enabled circular business models distinguishes between augmentation and automation, with AI respectively augmenting or automating parts of data-driven analysis and decision-making (Sjödin et al., 2023). Augmentation represents the more direct extension of the present framework: AI could support, rather than replace, managerial judgement by identifying emerging literature, market signals or regulatory developments that may warrant a review of the decision-

variable set, scenario assumptions or weights, provided that the organisational routines and AI capabilities required to operationalise such outputs are in place (Madanaguli et al., 2024). Such an extension would itself depend on the reliability, completeness and timeliness of the underlying information, an enabling condition for maintaining the DSS rather than a further decision variable within it. Future research could therefore examine how AI could support a more continuously informed review of the framework, while retaining final judgement with the decision-maker rather than transferring it to the AI.

These findings indicate that, in the case examined, the main constraint on a transition towards a more circular configuration is the financial architecture required for an OEM to retain control of the asset that contains the materials to be recovered, rather than the technical feasibility of recovering those materials itself. This is consistent with the wider CBM adoption literature, which identifies financial barriers as a prominent constraint (Hina et al., 2022), and the present case renders it traceable to two specific decision variables. The experts consulted for this thesis similarly emphasised financial viability as a key condition for adoption and, where financial considerations alone do not determine the outcome, identified binding regulatory requirements as another mechanism capable of influencing adoption. The case therefore suggests that the environmental benefits associated with greater material recovery do not automatically translate into the adoption of a more circular CBM. Their influence is mediated by the economic and institutional conditions under which the OEM makes the decision. The contribution of the DSS is to incorporate these conditions into the comparison of CBM alternatives and to show how the resulting recommendation changes when external conditions or the OEM's decision priorities change.

References

- Ahmed, A. A., Nazzal, M. A., Darras, B. M., & Deiab, I. M. (2022). A comprehensive multi-level circular economy assessment framework. *Sustainable Production and Consumption*, 32, 700-717.
- Assmann, I. R., Rosati, F., & Morioka, S. N. (2023). Determinants of circular business model adoption—A systematic literature review. *Business Strategy and the Environment*, 32(8), 6008-6028.
- Babbitt, C. W., Althaf, S., Rios, F. C., Bilec, M. M., & Graedel, T. E. (2021). The role of design in circular economy solutions for critical materials. *One Earth*, 4(3), 353-362.
- Baines, T. S., Lightfoot, H. W., Benedettini, O., & Kay, J. M. (2009). The servitization of manufacturing: A review of literature and reflection on future challenges. *Journal of Manufacturing Technology Management*, 20(5), 547-567.
- Baines, T. S., Lightfoot, H. W., Evans, S., Neely, A., Greenough, R., Peppard, J., ... & Wilson, H. (2007). State-of-the-art in product-service systems. *Proceedings of the Institution of Mechanical Engineers, Part B: Journal of Engineering Manufacture*, 221(10), 1543-1552.
- Baldassarre, B., & Carrara, S. (2025). Critical raw materials, circular economy, sustainable development: EU policy reflections for future research and innovation. *Resources, Conservation and Recycling*, 215, 108060.
- Beauson, J., Laurent, A., Rudolph, D. P., & Jensen, J. P. (2022). The complex end-of-life of wind turbine blades: A review of the European context. *Renewable and Sustainable Energy Reviews*, 155, 111847.
- Birt, L., Scott, S., Cavers, D., Campbell, C., & Walter, F. (2016). Member checking: a tool to enhance trustworthiness or merely a nod to validation? *Qualitative Health Research*, 26(13), 1802-1811.
- Bocken, N. M., De Pauw, I., Bakker, C., & Van Der Grinten, B. (2016). Product design and business model strategies for a circular economy. *Journal of Industrial and Production Engineering*, 33(5), 308-320.
- Carrara, S., Baldassarre, B., Jakimów, M., Kuzov, T., Mc Govern, L., Nohl, L., Ierides, M., & Christou, M. (2025). *Deep dive on critical raw materials for wind turbines in the EU* (EUR 40282; JRC141759). Publications Office of the European Union.
- Chinwego, C., Garcia, R., Sehar, D. M., MacGregor, E., Tate, S., Savage, K., ... & Mishra, B. (2025). Sufficiency-driven business models for rare earth recycling: integrating

- stakeholder collaboration and customer discovery for sustainable innovation. *Journal of Innovation and Entrepreneurship*, 14(1), 20.
- Cimprich, A., Young, S. B., Schrijvers, D., Ku, A. Y., Hagelüken, C., Christmann, P., ... & Hool, A. (2023). The role of industrial actors in the circular economy for critical raw materials: a framework with case studies across a range of industries. *Mineral Economics*, 36(2), 301-319.
- Cinelli, M., Kadziński, M., Gonzalez, M., & Słowiński, R. (2020). How to support the application of multiple criteria decision analysis? Let us start with a comprehensive taxonomy. *Omega*, 96, 102261.
- Couto, C. P., van de Ven, J. J., Yang, Y., & Abrahami, S. T. (2024). On the pre-treatment for recycling spent NdFeB permanent magnets: from disassembling, characterisation to de-coating. *Sustainable Materials and Technologies*, 41, e01041.
- Daou, A., Mallat, C., Chammass, G., Cerantola, N., Kayed, S., & Saliba, N. A. (2020). The Ecocanvas as a business model canvas for a circular economy. *Journal of Cleaner Production*, 258, 120938.
- Deveci, M., Delen, D., & Ding, W. (2024). AI-based decision support systems for sustainable business management under circular economy. *Annals of Operations Research*, 342(1), 1-3.
- Dresch, A., Lacerda, D. P., & Antunes Jr, J. A. V. (2015). *Design science research: A method for science and technology advancement*. Springer.
- Fleischmann, M., Bloemhof-Ruwaard, J. M., Dekker, R., Van der Laan, E., Van Nunen, J. A., & Van Wassenhove, L. N. (1997). Quantitative models for reverse logistics: A review. *European Journal of Operational Research*, 103(1), 1-17.
- Forman, E., & Peniwati, K. (1998). Aggregating individual judgments and priorities with the analytic hierarchy process. *European Journal of Operational Research*, 108(1), 165-169.
- Frederiksen, T. B., Pieroni, M. P. P., Pigosso, D. C., & McAloone, T. C. (2021). Strategic development of product-service systems (PSS) through archetype assessment. *Sustainability*, 13(5), 2592.
- Gaudenzi, B., Zsidisin, G. A., & Pellegrino, R. (2021). Measuring the financial effects of mitigating commodity price volatility in supply chains. *Supply Chain Management: An International Journal*, 26(1), 17-31.
- Gaustad, G., Krystofik, M., Bustamante, M., & Badami, K. (2018). Circular economy strategies for mitigating critical material supply issues. *Resources, Conservation and Recycling*, 135, 24-33.

- Geissdoerfer, M., Santa-Maria, T., Kirchherr, J., & Pelzeter, C. (2023). Drivers and barriers for circular business model innovation. *Business Strategy and the Environment*, 32(6), 3814-3832.
- Gioia, D. A., Corley, K. G., & Hamilton, A. L. (2013). Seeking qualitative rigor in inductive research: Notes on the Gioia methodology. *Organizational Research Methods*, 16(1), 15-31.
- Govindan, K., & Soleimani, H. (2017). A review of reverse logistics and closed-loop supply chains: a Journal of Cleaner Production focus. *Journal of Cleaner Production*, 142, 371-384.
- Govindan, K., Soleimani, H., & Kannan, D. (2015). Reverse logistics and closed-loop supply chain: A comprehensive review to explore the future. *European Journal of Operational Research*, 240(3), 603-626.
- Gregor, S., & Hevner, A. R. (2013). Positioning and presenting design science research for maximum impact. *MIS Quarterly*, 37(2), 337-355.
- Griffin, G., Gaustad, G., & Badami, K. (2019). A framework for firm-level critical material supply management and mitigation. *Resources Policy*, 60, 262-276.
- Grubic, T., & Jennions, I. (2018). Do outcome-based contracts exist? The investigation of power-by-the-hour and similar result-oriented cases. *International Journal of Production Economics*, 206, 209-219.
- Hevner, A. R. (2007). A three cycle view of design science research. *Scandinavian Journal of Information Systems*, 19(2), 4.
- Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75-105.
- Hidalgo-Crespo, J., Golinska-Dawson, P., & Riel, A. (2025). Product-as-a-service transition for original equipment manufacturers: Challenges, performance metrics, and design guidelines—The case of electrical and electronic equipment. *Cleaner Environmental Systems*, 19, 100330.
- Hina, M., Chauhan, C., Kaur, P., Kraus, S., & Dhir, A. (2022). Drivers and barriers of circular economy business models: Where we are now, and where we are heading. *Journal of Cleaner Production*, 333, 130049.
- Holsapple, C. W. (2008). DSS architecture and types. In *Handbook on decision support systems 1: Basic themes* (pp. 163-189). Berlin, Heidelberg: Springer Berlin Heidelberg.

- Hool, A., Helbig, C., & Wierink, G. (2024). Challenges and opportunities of the European Critical Raw Materials Act. *Mineral Economics*, 37(3), 661-668.
- Hörisch, J., Freeman, R. E., & Schaltegger, S. (2014). Applying stakeholder theory in sustainability management: Links, similarities, dissimilarities, and a conceptual framework. *Organization & Environment*, 27(4), 328-346.
- Hwang, C. L., & Yoon, K. (1981). *Multiple attribute decision making: methods and applications a state-of-the-art survey*. Springer-Verlag.
- Ibarra, D., Nieto Cerezo, O., & Mendoza, J. M. F. (2026). Circular business models and value chains for wind turbine blade end-of-life management: A roadmap for turning waste into value. *Renewable and Sustainable Energy Reviews*, 226, 116382.
- Ierides, M., Nohl, L., Alves Dias, P., Nohl, M., Blagoeva, D., & Magnani, N. (2025). *Substitution and reduction of critical and strategic raw materials in clean energy technologies* (EUR 40297; JRC142055). Publications Office of the European Union.
- Ishizaka, A., & Labib, A. (2011). Review of the main developments in the analytic hierarchy process. *Expert Systems with Applications*, 38(11), 14336-14345.
- Jin, T., Ding, Y., Guo, H., & Nalajala, N. (2012, July). Managing wind turbine reliability and maintenance via performance-based contract. In *2012 IEEE Power and Energy Society General Meeting* (pp. 1-6). IEEE.
- Kaliszewski, I., & Podkopaev, D. (2016). Simple additive weighting—A metamodel for multiple criteria decision analysis methods. *Expert Systems with Applications*, 54, 155-161.
- Kallio, H., Pietilä, A. M., Johnson, M., & Kangasniemi, M. (2016). Systematic methodological review: developing a framework for a qualitative semi-structured interview guide. *Journal of Advanced Nursing*, 72(12), 2954-2965.
- Kazancoglu, I., Sagnak, M., Kumar Mangla, S., & Kazancoglu, Y. (2021). Circular economy and the policy: A framework for improving the corporate environmental management in supply chains. *Business Strategy and the Environment*, 30(1), 590-608.
- Khor, K. S., & Udin, Z. M. (2013). Reverse logistics in Malaysia: Investigating the effect of green product design and resource commitment. *Resources, Conservation and Recycling*, 81, 71-80.
- Kirchherr, J., Piscicelli, L., Bour, R., Kostense-Smit, E., Muller, J., Huibrechtse-Truijens, A., & Hekkert, M. (2018). Barriers to the circular economy: Evidence from the European Union (EU). *Ecological Economics*, 150, 264-272.

- Kirchherr, J., Reike, D., & Hekkert, M. (2017). Conceptualizing the circular economy: An analysis of 114 definitions. *Resources, Conservation and Recycling*, 127, 221-232.
- Konietzko, J., Baldassarre, B., Brown, P., Bocken, N., & Hultink, E. J. (2020b). Circular business model experimentation: Demystifying assumptions. *Journal of Cleaner Production*, 277, 122596.
- Konietzko, J., Bocken, N., & Hultink, E. J. (2020a). Circular ecosystem innovation: An initial set of principles. *Journal of Cleaner Production*, 253, 119942.
- Kristensen, H. S., & Mosgaard, M. A. (2020). A review of micro level indicators for a circular economy—moving away from the three dimensions of sustainability? *Journal of Cleaner Production*, 243, 118531.
- Kristoffersen, E., Blomsma, F., Mikalef, P., & Li, J. (2020). The smart circular economy: A digital-enabled circular strategies framework for manufacturing companies. *Journal of Business Research*, 120, 241-261.
- Leader, A., Gaustad, G., & Babbitt, C. (2019). The effect of critical material prices on the competitiveness of clean energy technologies. *Materials for Renewable and Sustainable Energy*, 8(2), 8.
- Lewandowski, M. (2016). Designing the business models for circular economy—Towards the conceptual framework. *Sustainability*, 8(1), 43.
- Lüdeke-Freund, F., Gold, S., & Bocken, N. M. (2019). A review and typology of circular economy business model patterns. *Journal of Industrial Ecology*, 23(1), 36-61.
- Madanaguli, A., Sjödin, D., Parida, V., & Mikalef, P. (2024). Artificial intelligence capabilities for circular business models: Research synthesis and future agenda. *Technological Forecasting and Social Change*, 200, 123189.
- Marttunen, M., Mustajoki, J., Dufva, M., & Karjalainen, T. P. (2015). How to design and realize participation of stakeholders in MCDA processes? A framework for selecting an appropriate approach. *EURO Journal on Decision Processes*, 3(1), 187-214.
- Mendoza, J. M. F., Gallego-Schmid, A., Velenturf, A. P., Jensen, P. D., & Ibarra, D. (2022). Circular economy business models and technology management strategies in the wind industry: Sustainability potential, industrial challenges and opportunities. *Renewable and Sustainable Energy Reviews*, 163, 112523.
- Modgil, S., Gupta, S., Sivarajah, U., & Bhushan, B. (2021). Big data-enabled large-scale group decision making for circular economy: An emerging market context. *Technological Forecasting and Social Change*, 166, 120607.

- Mont, O. K. (2002). Clarifying the concept of product–service system. *Journal of Cleaner Production*, 10(3), 237-245.
- Muto, K., Kimita, K., Tanaka, H., Numata, E., Hosono, S., Izukura, S., & Shimomura, Y. (2016). A task management method for product service systems design. *Procedia CIRP*, 47, 537-542.
- Otterbach, N., & Fröhling, M. (2025). Raw material criticality assessments as decision-making support for corporate circular economy strategies—A systematic review. *Resources, Conservation and Recycling*, 219, 108299.
- Patton, M. Q. (2002). *Qualitative research and evaluation methods* (3rd ed.). Sage Publications.
- Pavel, C. C., Lacal-Arántegui, R., Marmier, A., Schüller, D., Tzimas, E., Buchert, M., ... & Blagoeva, D. (2017). Substitution strategies for reducing the use of rare earths in wind turbines. *Resources Policy*, 52, 349-357.
- Peffer, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, 24(3), 45-77.
- Petrova, V. (2023). Exploring the opportunities for sustainable management of critical raw materials in the circular economy. *The Eurasia Proceedings of Science Technology Engineering and Mathematics*, 26, 664-671.
- Pollard, J., Osmani, M., Cole, C., Grubnic, S., & Colwill, J. (2021). A circular economy business model innovation process for the electrical and electronic equipment sector. *Journal of Cleaner Production*, 305, 127211.
- Pollard, J., Osmani, M., Grubnic, S., Díaz, A. I., Grobe, K., Kaba, A., ... & Panchal, R. (2023). Implementing a circular economy business model canvas in the electrical and electronic manufacturing sector: A case study approach. *Sustainable Production and Consumption*, 36, 17-31.
- Ragonnaud, G. (2024). *Implementing the EU's Critical Raw Materials Act* (EPRS Briefing PE 766.253). European Parliamentary Research Service. [https://www.europarl.europa.eu/RegData/etudes/BRIE/2024/766253/EPRS_BRI\(2024\)766253_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/BRIE/2024/766253/EPRS_BRI(2024)766253_EN.pdf)
- Raut, S., Hossain, N. U. I., Kouhizadeh, M., & Fazio, S. A. (2025). Application of artificial intelligence in circular economy: A critical analysis of the current research. *Sustainable Futures*, 9, 100784.
- Regulation (EU) 2024/1252 of the European Parliament and of the Council of 11 April 2024 establishing a framework for ensuring a secure and sustainable supply of critical raw

- materials and amending Regulations (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1724 and (EU) 2019/1020. (2024). *Official Journal of the European Union*, L 2024/1252.
- Reim, W., Parida, V., & Örtqvist, D. (2015). Product–Service Systems (PSS) business models and tactics—a systematic literature review. *Journal of Cleaner Production*, 97, 61-75.
- Rizos, V., Righetti, E., & Kassab, A. (2024). Understanding the barriers to recycling critical raw materials for the energy transition: The case of rare earth permanent magnets. *Energy Reports*, 12, 1673-1682.
- Saaty, T. L. (1980). *The analytic hierarchy process: Planning, priority setting, resource allocation*. McGraw-Hill.
- Saaty, T. L. (1986). Absolute and relative measurement with the AHP. The most livable cities in the United States. *Socio-Economic Planning Sciences*, 20(6), 327-331.
- Saaty, T. L. (1990). How to make a decision: the analytic hierarchy process. *European Journal of Operational Research*, 48(1), 9-26.
- Saaty, T. L. (2008). Decision making with the analytic hierarchy process. *International Journal of Services Sciences*, 1(1), 83-98.
- Saidani, M., Yannou, B., Leroy, Y., Cluzel, F., & Kendall, A. (2019). A taxonomy of circular economy indicators. *Journal of Cleaner Production*, 207, 542-559.
- Schoemaker, P. J. (1995). Scenario planning: a tool for strategic thinking. *Sloan Management Review*, 36(2), 25-40.
- Siderius, T., & Zink, T. (2023). Markets and the future of the circular economy. *Circular Economy and Sustainability*, 3(3), 1569-1595.
- Sjödin, D., Parida, V., & Kohtamäki, M. (2023). Artificial intelligence enabling circular business model innovation in digital servitization: Conceptualizing dynamic capabilities, AI capacities, business models and effects. *Technological Forecasting and Social Change*, 197, 122903.
- Sprague Jr, R. H. (1980). A framework for the development of decision support systems. *MIS Quarterly*, 4(4), 1-26.
- Tang, J., Su, C., Wu, L., He, X., & Zhang, Y. (2026). Bargaining game analysis for performance-based maintenance contract decision-making of offshore wind turbines. *International Journal of Production Research*, 64(15), 6209-6222.

- Tranfield, D., Denyer, D., & Smart, P. (2003). Towards a methodology for developing evidence-informed management knowledge by means of systematic review. *British Journal of Management*, 14(3), 207-222.
- Tukker, A. (2004). Eight types of product–service system: eight ways to sustainability? Experiences from SusProNet. *Business Strategy and the Environment*, 13(4), 246-260.
- Tukker, A. (2015). Product services for a resource-efficient and circular economy—a review. *Journal of Cleaner Production*, 97, 76-91.
- Tyurkay, A., Kirkelund, G. M., & Lima, A. T. M. (2024). State-of-the-art circular economy practices for end-of-life wind turbine blades for use in the construction industry. *Sustainable Production and Consumption*, 47, 17-36.
- van Gaalen, J. M., & Slootweg, J. C. (2025). From critical raw materials to circular raw materials. *ChemSusChem*, 18(2), e202401170.
- von Rennenberg, T., Yu, Y., Fraccascia, L., & Yazan, D. M. (2024). Circularity Reinforcement of Critical Raw Materials in Europe: A Case of Niobium. *Circular Economy and Sustainability*, 4(4), 2449-2472.
- Wang, J., Zhao, X., & Guo, X. (2019). Optimizing wind turbine's maintenance policies under performance-based contract. *Renewable Energy*, 135, 626-634.
- Wiles, R., Crow, G., Heath, S., & Charles, V. (2008). The management of confidentiality and anonymity in social research. *International Journal of Social Research Methodology*, 11(5), 417-428.
- Yang, M., Smart, P., Kumar, M., Jolly, M., & Evans, S. (2018). Product-service systems business models for circular supply chains. *Production Planning & Control*, 29(6), 498-508.

Appendix A

A. Interview Metadata and Logistics (interviewer-only)

Date	
Interviewee	
Role / Organisation	
Interview format	Online video call (Microsoft Teams) with screen sharing, or in-person session with the guiding document provided in printed form
Planned duration	≈ 60 minutes
Recording	Yes / No (to be confirmed at the start of the call)
Consent obtained	Yes / No
Identification preference	Named (name and role) / Anonymised
Interviewer	Matilde Delfino

Tab. A.1 - Interview metadata, completed once per session.

B. Opening Script: Purpose, Consent and Confidentiality

“Thank you very much for taking the time to speak with me today. This research develops a Decision Support System (DSS) to help an Original Equipment Manufacturer (OEM) choose, in a structured and transparent way, among three CBM archetypes for products that rely heavily on critical raw materials, under conditions of price volatility in those materials. The OEM is treated as the focal decision-maker, given its central position in the value chain and its influence over product design and end-of-life recovery. The framework is demonstrated on a case in the offshore wind sector, focusing on permanent-magnet generators and the recovery of the critical raw materials they contain. This interview has two parts. In the first part, I’ll ask a few open questions drawing on your professional experience and what you have observed in the sector; there are no right or wrong answers here. In the second part, I’ll ask you to compare a set of decision criteria two at a time, adopting the perspective of the OEM; I’ll explain exactly how that works when we get there. Everything you say will be treated confidentially and used for academic research purposes only. Where your input is referenced in the thesis, I will ask for your confirmation of how you’d like to be identified: by name and role, or anonymised. Would you be comfortable if I recorded this call for transcription purposes only? And do you have any questions before we start?”

C. Warm-up Questions: Interviewee Background

Open questions. There are no right or wrong answers; the interviewer is simply interested in the interviewee’s background and what they have observed in the sector.

1. Could you briefly describe your current role, and how it relates to offshore wind and to OEMs (turbine/generator manufacturers) specifically, if at all?
2. How many years of experience do you have in the renewable energy / offshore wind sector?

3. Could you briefly describe your professional background and how it connects to offshore wind and to OEMs specifically?
4. Have you worked directly with, or within, an OEM (turbine/generator manufacturer) at any point? If so, in what capacity?
5. Are you familiar with servitisation or “as-a-service” business models in this sector (e.g. leasing, performance-based contracts, pay-per-output)?

D. Framework Review: Clarity, Relevance and Completeness

Open questions. This part reviews the framework developed in Section 3.2 before the structured AHP exercise. The decision variables were defined before the interviews; this review served as a qualitative check of their clarity, relevance and completeness, and was not intended to generate new variables.

1. Looking at the six macro-categories of the framework, are they clear and readily understandable from the perspective of the OEM’s decision problem?
2. Do the decision variables within these categories appear relevant to the OEM’s decision when choosing among the three CBM archetypes?
3. Considered together, do the six categories and their decision variables seem sufficiently comprehensive, or does any category appear to overlap with another?
4. Is there any relevant factor or decision variable that you think is missing from the framework?

E. Transition Script: Framing the OEM Decision-Maker Perspective

Read aloud before the first comparison. This fixes the vantage point from which every subsequent importance judgement is made; it states only the perspective to be adopted.

“From this point on, please answer as if you were the OEM itself, the manufacturer of the offshore wind turbine, choosing among the three CBMs just introduced. What matters is not your own institutional viewpoint, but how important each criterion is to the OEM’s decision, regardless of which model you might personally expect to score best on it.”

Given your experience in the sector rather than being an OEM employee yourself, do you feel comfortable answering from the OEM’s own decision-making perspective on this basis? If the interviewee expresses reservations, note them explicitly for the thesis’s Limitations section and proceed only if they remain willing to give their best-informed judgement.

F. Structured AHP Comparisons

For every comparison, the interviewee indicates which of the two elements matters more to the OEM's decision (or whether they matter equally), and how much more, using the AHP standard 1-9 comparison scale (Saaty, 1980).

Value	Meaning
1	Equal importance: the two elements contribute equally to the decision.
3	Moderate importance of one element over the other.
5	Strong importance of one element over the other.
7	Very strong / demonstrated importance of one element over the other.
9	Extreme importance of one element over the other.
2, 4, 6, 8	Intermediate values, for judgements that fall between two of the levels above.

Tab. A.2 - Saaty Fundamental Scale (1-9).

Before starting the actual comparisons, the interviewer provides a short example unrelated to the research topic to illustrate how the 1-9 scale is used.

Note: the “thinking as the OEM” framing given in Section E, together with a reminder of the meaning of each category or variable involved (using the definitions in Section 3.2 / Chapter 4 of the thesis, without disclosing benefit/cost classification or literature-based performance ratings), is repeated by the interviewer immediately before every comparison matrix below.

“Thinking as the OEM ... which of these two matters more to that decision, A or B? And how much more important, on the scale above?”

F.1 Level 1: Comparing the Six Macro-Categories (Market, Economic/Business Case, Technical/Product, Organisational/Capabilities, Value Chain/Ecosystem, Regulatory/Institutional)

Considering the overall goal of selecting the most suitable CBM archetype under CRM price volatility, and thinking as the OEM:

#	Element A	Element B	Winner (A / B / Equal)	Intensity (1-9)
1	Market (MKT)	Economic / Business Case (ECO)		
2	Market (MKT)	Technical / Product (TEC)		
3	Market (MKT)	Organisational / Capabilities (ORG)		
4	Market (MKT)	Value Chain / Ecosystem (VCE)		
5	Market (MKT)	Regulatory / Institutional (REG)		
6	Economic / Business Case (ECO)	Technical / Product (TEC)		

#	Element A	Element B	Winner (A / B / Equal)	Intensity (1-9)
7	Economic / Business Case (ECO)	Organisational / Capabilities (ORG)		
8	Economic / Business Case (ECO)	Value Chain / Ecosystem (VCE)		
9	Economic / Business Case (ECO)	Regulatory / Institutional (REG)		
10	Technical / Product (TEC)	Organisational / Capabilities (ORG)		
11	Technical / Product (TEC)	Value Chain / Ecosystem (VCE)		
12	Technical / Product (TEC)	Regulatory / Institutional (REG)		
13	Organisational / Capabilities (ORG)	Value Chain / Ecosystem (VCE)		
14	Organisational / Capabilities (ORG)	Regulatory / Institutional (REG)		
15	Value Chain / Ecosystem (VCE)	Regulatory / Institutional (REG)		

Tab. A.3 - Level 1 pairwise comparisons (15 pairs, fixed order).

F.2 Level 2: Comparing the Variables within Each Category

Repeat the same procedure independently within each category, comparing only the variables that belong to it, after reminding the interviewee of the category framing.

Comparing the Variables within: Market

#	Element A	Element B	Winner (A / B / Equal)	Intensity (1-9)
1	Customer Acceptance of the CBM's Ownership/Access Arrangement	Customer Asset-Care and Moral-Hazard Exposure		
2	Customer Acceptance of the CBM's Ownership/Access Arrangement	Demand/Usage Forecasting Dependency		
3	Customer Asset-Care and Moral-Hazard Exposure	Demand/Usage Forecasting Dependency		

Tab. A.4 - Within-category pairwise comparisons: Market.

Comparing the Variables within: Economic / Business Case

#	Element A	Element B	Winner (A / B / Equal)	Intensity (1-9)
1	Provider Investment/Capital Requirement	Operational Service-Delivery Cost Burden		
2	Provider Investment/Capital Requirement	Value Creation and Capture Potential		
3	Provider Investment/Capital Requirement	CRM Price Volatility Exposure		
4	Operational Service-Delivery Cost Burden	Value Creation and Capture Potential		
5	Operational Service-Delivery Cost Burden	CRM Price Volatility Exposure		
6	Value Creation and Capture Potential	CRM Price Volatility Exposure		

Tab. A.5 - Within-category pairwise comparisons: Economic / Business Case.

Comparing the Variables within: Technical / Product

#	Element A	Element B	Winner (A / B / Equal)	Intensity (1-9)
1	Technological Infrastructure and Complexity Requirement	Product Reliability and Durability		
2	Technological Infrastructure and Complexity Requirement	Technology and Material Substitution Risk over the Contract Horizon		
3	Product Reliability and Durability	Technology and Material Substitution Risk over the Contract Horizon		

Tab. A.6 - Within-category pairwise comparisons: Technical / Product.

Comparing the Variables within: Organisational / Capabilities

#	Element A	Element B	Winner (A / B / Equal)	Intensity (1-9)
1	Service and Maintenance Capability	Organisational Restructuring Requirement		

Tab. A.7 - Within-category pairwise comparisons: Organisational / Capabilities.

Comparing the Variables within: Value Chain / Ecosystem

#	Element A	Element B	Winner (A / B / Equal)	Intensity (1-9)
1	CRM Recovery and Asset Return Effectiveness	Lifecycle Data Accessibility and Control		
2	CRM Recovery and Asset Return Effectiveness	Ecosystem Coordination Complexity		
3	Lifecycle Data Accessibility and Control	Ecosystem Coordination Complexity		

Tab. A.8 - Within-category pairwise comparisons: Value Chain / Ecosystem.

Comparing the Variables within: Regulatory / Institutional

#	Element A	Element B	Winner (A / B / Equal)	Intensity (1-9)
1	Regulatory Traceability Compliance Capability	Regulatory and Financial Incentive Availability		
2	Regulatory Traceability Compliance Capability	Exposure to Regulatory Change over the Contract Horizon		
3	Regulatory and Financial Incentive Availability	Exposure to Regulatory Change over the Contract Horizon		

Tab. A.9 - Within-category pairwise comparisons: Regulatory / Institutional.

G. Closing Questions

1. Reflecting on the comparisons we have just gone through, which of these considerations do you think tend to matter most in practice for this kind of decision?
2. Is there anything about this decision, from the OEM's perspective, that you feel wasn't captured by these categories or variables?

Appendix B

This appendix reports the Python implementation of the Decision Support System described in Section 4.5, used to compute every consistency ratio, weight and scenario score reported in Chapter 5. Section B.1 gives a short overview of the software and of the input workbook it reads, Section B.2 reproduces the source code in full, and Section B.3 reports the output obtained on the case-study data.

B.1 Overview

The pipeline relies on two external libraries, NumPy for the eigenvector computation and openpyxl for reading the input workbook and writing the output report and was run from a Jupyter notebook on the case-study data; saved as a single file, the same code can equally be executed as a stand-alone script. The input workbook contains a legend of the hierarchy codes, one sheet per expert with the seven pairwise comparison matrices, and a Performance_Matrix sheet with the normalised scores of Tabs. 4.13 and 4.14 for the baseline, Low and High conditions; only the upper triangle of each comparison matrix is entered, the lower triangle being its reciprocal by construction. These sheets are shown in Figs. B.1 to B.6, using the judgements after the revision of the matrices described in Section 5.2.1.

	A	B	C	D
1	Code	Level	Macro-category	Full name
2	MKT	Category	Market	
3	MKT1	Variable	Market	Customer Acceptance of the CBM's Ownership/Access Arrangement
4	MKT2	Variable	Market	Customer Asset-Care and Moral-Hazard Exposure
5	MKT3	Variable	Market	Demand/Usage Forecasting Dependency
6	ECO	Category	Economic/Business Case	
7	ECO1	Variable	Economic/Business Case	Provider Investment/Capital Requirement
8	ECO2	Variable	Economic/Business Case	Operational Service-Delivery Cost Burden
9	ECO3	Variable	Economic/Business Case	Value Creation and Capture Potential
10	ECO4	Variable	Economic/Business Case	CRM Price Volatility Exposure
11	TEC	Category	Technical/Product	
12	TEC1	Variable	Technical/Product	Technological Infrastructure and Complexity Requirement
13	TEC2	Variable	Technical/Product	Product Reliability and Durability
14	TEC3	Variable	Technical/Product	Technology and Material Substitution Risk over the Contract Horizon
15	ORG	Category	Organisational/Capabilities	
16	ORG1	Variable	Organisational/Capabilities	Service and Maintenance Capability
17	ORG2	Variable	Organisational/Capabilities	Organisational Restructuring Requirement
18	VCE	Category	Value Chain/Ecosystem	
19	VCE1	Variable	Value Chain/Ecosystem	CRM Recovery and Asset Return Effectiveness
20	VCE2	Variable	Value Chain/Ecosystem	Lifecycle Data Accessibility and Control
21	VCE3	Variable	Value Chain/Ecosystem	Ecosystem Coordination Complexity
22	REG	Category	Regulatory/Institutional	
23	REG1	Variable	Regulatory/Institutional	Regulatory Traceability Compliance Capability
24	REG2	Variable	Regulatory/Institutional	Regulatory and Financial Incentive Availability
25	REG3	Variable	Regulatory/Institutional	Exposure to Regulatory Change over the Contract Horizon

Fig. B.1 - Legend sheet of the input workbook.

	A	B	C	D	E	F	G	H
1	Matrix 1 - Comparison of the 6 Macro-categories (Expert 1)							
2		MKT	ECO	TEC	ORG	VCE	REG	
3	MKT	1	1/9	1/7	1	1/5	1/3	
4	ECO	9	1	7	5	3	1	
5	TEC	7	1/7	1	5	1	1	
6	ORG	1	1/5	1/5	1	1/7	1/5	
7	VCE	5	1/3	1	7	1	1	
8	REG	3	1	1	5	1	1	
9								
10								
11	Matrix 2 - Market (MKT) - comparison of its 3 variables							
12		MKT1	MKT2	MKT3				
13	MKT1	1	3	1				
14	MKT2	1/3	1	1/3				
15	MKT3	1	3	1				
16								
17								
18	Matrix 3 - Economic/Business Case (ECO) - comparison of its 4 variables							
19		ECO1	ECO2	ECO3	ECO4			
20	ECO1	1	1	5	1			
21	ECO2	1	1	2	1			
22	ECO3	1/5	1/2	1	1/5			
23	ECO4	1	1	5	1			
24								
25								
26	Matrix 4 - Technical/Product (TEC) - comparison of its 3 variables							
27		TEC1	TEC2	TEC3				
28	TEC1	1	1/5	1/7				
29	TEC2	5	1	1				
30	TEC3	7	1	1				
31								
32								
33	Matrix 5 - Organisational/Capabilities (ORG) - comparison of its 2 variables							
34		ORG1	ORG2					
35	ORG1	1	5					
36	ORG2	1/5	1					
37								
38								
39	Matrix 6 - Value Chain/Ecosystem (VCE) - comparison of its 3 variables							
40		VCE1	VCE2	VCE3				
41	VCE1	1	1	5				
42	VCE2	1	1	7				
43	VCE3	1/5	1/7	1				
44								
45								
46	Matrix 7 - Regulatory/Institutional (REG) - comparison of its 3 variables							
47		REG1	REG2	REG3				
48	REG1	1	1/3	1/4				
49	REG2	3	1	1				
50	REG3	4	1	1				

Fig. B.2 - Expert_1 sheet: pairwise comparison matrices of expert E1.

	A	B	C	D	E	F	G	H
1	Matrix 1 - Comparison of the 6 Macro-categories (Expert 2)							
2		MKT	ECO	TEC	ORG	VCE	REG	
3	MKT	1	1/5	3	3	5	5	
4	ECO	5	1	5	3	5	5	
5	TEC	1/3	1/5	1	1	3	2	
6	ORG	1/3	1/3	1	1	1	2	
7	VCE	1/5	1/5	1/3	1	1	2	
8	REG	1/5	1/5	1/2	1/2	1/2	1	
9								
10								
11	Matrix 2 - Market (MKT) - comparison of its 3 variables							
12		MKT1	MKT2	MKT3				
13	MKT1	1	1	3				
14	MKT2	1	1	2				
15	MKT3	1/3	1/2	1				
16								
17								
18	Matrix 3 - Economic/Business Case (ECO) - comparison of its 4 variables							
19		ECO1	ECO2	ECO3	ECO4			
20	ECO1	1	1/2	1/3	1/2			
21	ECO2	2	1	1/3	1			
22	ECO3	3	3	1	3			
23	ECO4	2	1	1/3	1			
24								
25								
26	Matrix 4 - Technical/Product (TEC) - comparison of its 3 variables							
27		TEC1	TEC2	TEC3				
28	TEC1	1	1/3	5				
29	TEC2	3	1	7				
30	TEC3	1/5	1/7	1				
31								
32								
33	Matrix 5 - Organisational/Capabilities (ORG) - comparison of its 2 variables							
34		ORG1	ORG2					
35	ORG1	1	3					
36	ORG2	1/3	1					
37								
38								
39	Matrix 6 - Value Chain/Ecosystem (VCE) - comparison of its 3 variables							
40		VCE1	VCE2	VCE3				
41	VCE1	1	1/5	1				
42	VCE2	5	1	3				
43	VCE3	1	1/3	1				
44								
45								
46	Matrix 7 - Regulatory/Institutional (REG) - comparison of its 3 variables							
47		REG1	REG2	REG3				
48	REG1	1	3	1/3				
49	REG2	1/3	1	1/5				
50	REG3	3	5	1				

Fig. B.3 - Expert_2 sheet: pairwise comparison matrices of expert E2.

	A	B	C	D	E	F	G	H
1	Matrix 1 - Comparison of the 6 Macro-categories (Expert 3)							
2		MKT	ECO	TEC	ORG	VCE	REG	
3	MKT	1	1/7	3	1	5	1/5	
4	ECO	7	1	9	7	9	3	
5	TEC	1/3	1/9	1	1/3	3	1/5	
6	ORG	1	1/7	3	1	3	1/3	
7	VCE	1/5	1/9	1/3	1/3	1	1/5	
8	REG	5	1/3	5	3	5	1	
9								
10								
11	Matrix 2 - Market (MKT) - comparison of its 3 variables							
12		MKT1	MKT2	MKT3				
13	MKT1	1	7	9				
14	MKT2	1/7	1	3				
15	MKT3	1/9	1/3	1				
16								
17								
18	Matrix 3 - Economic/Business Case (ECO) - comparison of its 4 variables							
19		ECO1	ECO2	ECO3	ECO4			
20	ECO1	1	1	1	8			
21	ECO2	1	1	1	9			
22	ECO3	1	1	1	9			
23	ECO4	1/8	1/9	1/9	1			
24								
25								
26	Matrix 4 - Technical/Product (TEC) - comparison of its 3 variables							
27		TEC1	TEC2	TEC3				
28	TEC1	1	1	3				
29	TEC2	1	1	3				
30	TEC3	1/3	1/3	1				
31								
32								
33	Matrix 5 - Organisational/Capabilities (ORG) - comparison of its 2 variables							
34		ORG1	ORG2					
35	ORG1	1	1/7					
36	ORG2	7	1					
37								
38								
39	Matrix 6 - Value Chain/Ecosystem (VCE) - comparison of its 3 variables							
40		VCE1	VCE2	VCE3				
41	VCE1	1	1/5	1/3				
42	VCE2	5	1	3				
43	VCE3	3	1/3	1				
44								
45								
46	Matrix 7 - Regulatory/Institutional (REG) - comparison of its 3 variables							
47		REG1	REG2	REG3				
48	REG1	1	1/7	1/5				
49	REG2	7	1	3				
50	REG3	5	1/3	1				

Fig. B.4 - Expert_3 sheet: pairwise comparison matrices of expert E3.

	A	B	C	D	E	F	G	H
1	Matrix 1 - Comparison of the 6 Macro-categories (Expert 4)							
2		MKT	ECO	TEC	ORG	VCE	REG	
3	MKT	1	1/3	7	5	9	3	
4	ECO	3	1	7	5	9	3	
5	TEC	1/7	1/7	1	1/3	3	1/5	
6	ORG	1/5	1/5	3	1	5	1/3	
7	VCE	1/9	1/9	1/3	1/5	1	1/7	
8	REG	1/3	1/3	5	3	7	1	
9								
10								
11	Matrix 2 - Market (MKT) - comparison of its 3 variables							
12		MKT1	MKT2	MKT3				
13	MKT1	1	3	5				
14	MKT2	1/3	1	3				
15	MKT3	1/5	1/3	1				
16								
17								
18	Matrix 3 - Economic/Business Case (ECO) - comparison of its 4 variables							
19		ECO1	ECO2	ECO3	ECO4			
20	ECO1	1	1	9	7			
21	ECO2	1	1	7	3			
22	ECO3	1/9	1/7	1	1/3			
23	ECO4	1/7	1/3	3	1			
24								
25								
26	Matrix 4 - Technical/Product (TEC) - comparison of its 3 variables							
27		TEC1	TEC2	TEC3				
28	TEC1	1	1/7	1/3				
29	TEC2	7	1	5				
30	TEC3	3	1/5	1				
31								
32								
33	Matrix 5 - Organisational/Capabilities (ORG) - comparison of its 2 variables							
34		ORG1	ORG2					
35	ORG1	1	7					
36	ORG2	1/7	1					
37								
38								
39	Matrix 6 - Value Chain/Ecosystem (VCE) - comparison of its 3 variables							
40		VCE1	VCE2	VCE3				
41	VCE1	1	1/7	1/3				
42	VCE2	7	1	3				
43	VCE3	3	1/3	1				
44								
45								
46	Matrix 7 - Regulatory/Institutional (REG) - comparison of its 3 variables							
47		REG1	REG2	REG3				
48	REG1	1	1/3	3				
49	REG2	3	1	5				
50	REG3	1/3	1/5	1				

Fig. B.5 - Expert_4 sheet: pairwise comparison matrices of expert E4.

A	B	C	D	E	F	G	H	I	J	K	L	M	N
Code	Macro-category	Decision variable	Polarity	Scenario- dependent	Product (Baseline)	Use (Baseline)	Result (Baseline)	Product (Low)	Use (Low)	Result (Low)	Product (High)	Use (High)	Result (High)
1	Code												
2	MKT1	Market	Customer Acceptance of the CBM's Ownership/Access Arrangement	Benefit	Yes	1.00	0.67	0.33	0.67	0.33	0.67	1.00	0.67
3	MKT2	Market	Customer Asset-Care and Moral-Hazard Exposure	Cost	No	1.00	0.33	0.67	0.33	0.67	1.00	0.33	0.67
4	MKT3	Market	Demand/Usage Forecasting Dependency	Cost	No	1.00	0.67	0.33	0.33	0.67	1.00	0.67	0.33
5	ECO1	Economic/Business Case	Provider Investment/Capital Requirement	Cost	No	1.00	0.67	0.33	0.67	0.33	1.00	0.67	0.33
6	ECO2	Economic/Business Case	Operational Service-Delivery Cost Burden	Cost	Yes	1.00	0.67	0.33	1.00	0.67	0.67	0.33	0.33
7	ECO3	Economic/Business Case	Value Creation and Capture Potential	Benefit	Yes	0.67	1.00	1.00	0.67	0.67	0.33	1.00	1.00
8	ECO4	Economic/Business Case	CRM Price Volatility Exposure	Cost	Yes	0.33	0.67	1.00	1.00	0.67	0.33	0.67	1.00
9	TEC1	Technical/Product	Technological Infrastructure and Complexity Requirement	Cost	No	1.00	0.67	0.33	0.67	0.33	1.00	0.67	0.33
10	TEC2	Technical/Product	Product Reliability and Durability	Benefit	No	0.33	1.00	0.33	1.00	1.00	0.33	1.00	1.00
11	TEC3	Technical/Product	Technology and Material Substitution Risk over the Contract Horizon	Cost	Yes	1.00	0.33	0.67	0.67	1.00	0.67	0.33	0.33
12	ORG1	Organisational/Capabilities	Service and Maintenance Capability	Cost	Yes	1.00	0.67	0.33	0.67	0.33	0.67	0.33	0.33
13	ORG2	Organisational/Capabilities	Organisational Restructuring Requirement	Cost	No	1.00	0.67	0.33	0.67	0.33	1.00	0.67	0.33
14	VCE1	Value Chain/Ecosystem	CRM Recovery and Asset Return Effectiveness	Benefit	No	0.33	0.67	1.00	0.67	1.00	0.33	0.67	1.00
15	VCE2	Value Chain/Ecosystem	Lifecycle Data Accessibility and Control	Benefit	No	0.33	0.67	1.00	0.67	1.00	0.33	0.67	1.00
16	VCE3	Value Chain/Ecosystem	Ecosystem Coordination Complexity	Cost	No	0.33	0.67	0.33	0.67	0.33	0.33	0.67	0.33
17	REG1	Regulatory/Institutional	Regulatory Traceability Compliance Capability	Benefit	No	0.33	0.67	1.00	0.67	1.00	0.33	0.67	1.00
18	REG2	Regulatory/Institutional	Regulatory and Financial Incentive Availability	Benefit	No	0.33	1.00	0.33	1.00	1.00	0.33	1.00	1.00
19	REG3	Regulatory/Institutional	Exposure to Regulatory Change over the Contract Horizon	Cost	No	1.00	0.33	0.33	0.33	0.33	1.00	0.33	0.33

Fig. B.6 - Performance_Matrix sheet: normalised scores of the three archetypes under the Baseline, Low and High conditions.

B.2 Source Code Listing

Listing B.1 reproduces the code used for the analysis of Chapter 5, including its original comments; a small number of variable- and file-name inconsistencies present in the code as originally executed have been corrected here for clarity, with no effect on the computations or results reported in Chapter 5. Lines longer than the width of the page are wrapped and continue on the following line without a line number.

```
1  """
2  AHP Decision Support System - CBM Offshore Wind Thesis
3  Implements Chapter 3 methodology (Steps 0-6):
4  0: read and validate the Excel data file
5  1: per-expert, per-matrix priority vector (eigenvector) + individual CR (eq. 3.7-3.13)
6  2: aggregation of individual judgements (AIJ) via geometric mean (eq. 3.9)
7  3: group-level CR check on the aggregated matrices (Sec. 3.3.3)
8  4: local and global weights,  $W_i = W_c \times w_i/c$  (eq. 3.10)
9  5: read the fixed Performance Matrix (Product/Use/Result x Baseline/Low/High)
10 6: final scenario scores and ranking (eq. 3.17-3.21)
11
12 Usage:
13     the input and output filenames are set in the last four lines of this file; run: python
14     ahp_analysis.py
15 """
16 import numpy as np
17 import openpyxl
18
19 CR_THRESHOLD = 0.10
20
21 # Saaty (1980) Random Index, Tab. 3.2 in the thesis, indexed by matrix size n
22 RANDOM_INDEX = {1: 0.00, 2: 0.00, 3: 0.58, 4: 0.90, 5: 1.12,
23                 6: 1.24, 7: 1.32, 8: 1.41, 9: 1.45, 10: 1.49}
24
25 # Decision hierarchy - must mirror the structure of AHP_Data_Template.xlsx
26 HIERARCHY = {
27     "MKT": ("Market", ["MKT1", "MKT2", "MKT3"]),
28     "ECO": ("Economic/Business Case", ["ECO1", "ECO2", "ECO3", "ECO4"]),
29     "TEC": ("Technical/Product", ["TEC1", "TEC2", "TEC3"]),
30     "ORG": ("Organisational/Capabilities", ["ORG1", "ORG2"]),
31     "VCE": ("Value Chain/Ecosystem", ["VCE1", "VCE2", "VCE3"]),
32     "REG": ("Regulatory/Institutional", ["REG1", "REG2", "REG3"]),
33 }
34 CAT_CODES = list(HIERARCHY.keys())
35 ARCHETYPES = ["Product", "Use", "Result"]
36 SCENARIOS = ["Baseline", "Low", "High"]
37
38
39 # -----
40 # STEP 0 - Locate and read the pairwise-comparison matrices from a sheet
41 # -----
42 def matrix_layout():
43     """Returns (key, codes, header_row) for each of the 7 matrices, matching
44     the row layout used in the template."""
45     layout = []
46     start_row = 1
47     matrices = [{"CAT": "CAT", CAT_CODES}] + [(code, vars_) for code, (_, vars_) in
48 HIERARCHY.items()]
49     for key, codes in matrices:
50         n = len(codes)
51         header_row = start_row + 1
52         layout.append((key, codes, header_row))
53         start_row = header_row + 1 + n + 2
54     return layout
55
56 def read_matrix(ws, codes, header_row):
57     """Reads an n x n reciprocal matrix from its upper triangle; raises if a
58     required judgement cell is missing."""
```

```

59     n = len(codes)
60     A = np.ones((n, n))
61     for i in range(n):
62         for j in range(n):
63             if i == j:
64                 continue
65             row = header_row + 1 + i
66             col = 2 + j
67             val = ws.cell(row=row, column=col).value
68             if val is None or val == "":
69                 if j > i: # upper triangle must be filled by the user
70                     raise ValueError(
71                         f"[{ws.title}] Missing judgement at matrix starting row "
72                         f"{header_row}, cell for '{codes[i]}' vs '{codes[j]}' "
73                         f"(row {row}, col {openpyxl.utils.get_column_letter(col)})."
74                     )
75                 continue
76             A[i, j] = float(val)
77     return A
78
79
80 def read_expert_sheet(ws):
81     """Returns dict: matrix_key -> (codes, matrix) for the 7 matrices on one expert
82     sheet."""
83     return {key: (codes, read_matrix(ws, codes, header_row))
84             for key, codes, header_row in matrix_layout()}
85
86 # -----
87 # AHP core: priority vector (principal eigenvector) + Consistency Ratio
88 # -----
89 def priority_vector_and_cr(A):
90     n = A.shape[0]
91     eigvals, eigvecs = np.linalg.eig(A)
92     idx = np.argmax(eigvals.real)
93     lam_max = eigvals.real[idx]
94     w = np.abs(eigvecs[:, idx].real)
95     w = w / w.sum()
96     if n <= 2:
97         return w, lam_max, 0.0 # CR undefined/zero for n<=2 by construction
98     CI = (lam_max - n) / (n - 1)
99     RI = RANDOM_INDEX[n]
100     CR = CI / RI if RI > 0 else 0.0
101     CR = max(CR, 0.0) # a perfectly consistent matrix can give a tiny negative CR from fp
102     return w, lam_max, CR
103
104
105 def geometric_mean_aggregate(matrices):
106     """AIJ aggregation: cell-wise geometric mean across experts (eq. 3.9)."""
107     stacked = np.stack(matrices, axis=0)
108     return np.exp(np.mean(np.log(stacked), axis=0))
109
110
111 # -----
112 # STEP 5 - Performance matrix (fixed, literature-based, Chapter 4)
113 # -----
114 def read_performance_matrix(ws):
115     perf = {}
116     headers = {}
117     for c in range(1, ws.max_column + 1):
118         h = ws.cell(row=1, column=c).value
119         if h:
120             headers[h.strip()] = c
121     col_map = {
122         "Baseline": {"Product": headers["Product (Baseline)"], "Use": headers["Use
123         (Baseline)"],
124                     "Result": headers["Result (Baseline)"]},
125         "Low": {"Product": headers["Product (Low)"], "Use": headers["Use (Low)"],
126               "Result": headers["Result (Low)"]},
127         "High": {"Product": headers["Product (High)"], "Use": headers["Use (High)"],

```

```

127         "Result": headers["Result (High)"]},
128     }
129     code_col = headers["Code"]
130     r = 2
131     while ws.cell(row=r, column=code_col).value:
132         code = ws.cell(row=r, column=code_col).value.strip()
133         perf[code] = {}
134         for scen in SCENARIOS:
135             perf[code][scen] = {
136                 arch: float(ws.cell(row=r, column=col_map[scen][arch]).value)
137                 for arch in ARCHETYPES
138             }
139             r += 1
140     return perf
141
142
143 # -----
144 # Main pipeline
145 # -----
146 def find_expert_sheets(wb):
147     """Finds all 'Expert_N' sheets, sorted numerically (supports any number of experts)."""
148     import re
149     sheets = []
150     for name in wb.sheetnames:
151         m = re.fullmatch(r"Expert_(\d+)", name)
152         if m:
153             sheets.append((int(m.group(1)), name))
154     sheets.sort()
155     if not sheets:
156         raise ValueError("No sheet matching 'Expert_N' found in workbook.")
157     return [name for _, name in sheets]
158
159
160 def run(input_path):
161     wb = openpyxl.load_workbook(input_path, data_only=True)
162     expert_sheets = find_expert_sheets(wb)
163
164     # STEP 0/1: read + per-expert priority vectors & individual CR
165     experts_data = {} # expert -> {matrix_key: (codes, matrix)}
166     individual_report = []
167     for s in expert_sheets:
168         ws = wb[s]
169         experts_data[s] = read_expert_sheet(ws)
170         for key, (codes, A) in experts_data[s].items():
171             w, lam_max, CR = priority_vector_and_cr(A)
172             flag = "OK" if CR <= CR_THRESHOLD else "REVIEW (CR > 0.10)"
173             individual_report.append({
174                 "expert": s, "matrix": key, "n": len(codes),
175                 "lambda_max": lam_max, "CR": CR, "flag": flag,
176                 "weights": dict(zip(codes, w)),
177             })
178
179     # STEP 2: AIJ aggregation (geometric mean across experts)
180     aggregated = {} # matrix_key -> (codes, aggregated_matrix)
181     for key, codes, header_row in matrix_layout():
182         mats = [experts_data[s][key][1] for s in expert_sheets]
183         aggregated[key] = (codes, geometric_mean_aggregate(mats))
184
185     # STEP 3: group-level CR on aggregated matrices
186     group_report = []
187     group_weights = {} # matrix_key -> {code: weight}
188     for key, (codes, A) in aggregated.items():
189         w, lam_max, CR = priority_vector_and_cr(A)
190         flag = "OK" if CR <= CR_THRESHOLD else "REVIEW (CR > 0.10)"
191         group_report.append({"matrix": key, "n": len(codes), "lambda_max": lam_max,
192                             "CR": CR, "flag": flag})
193         group_weights[key] = dict(zip(codes, w))
194
195     overall_cr_ok = all(r["flag"] == "OK" for r in group_report)
196
197     # STEP 4: Local + global weights,  $W_i = W_c \times w_i/c$ 

```

```

198 Wc = group_weights["CAT"] # category weights
199 global_weights = {}
200 variable_meta = {}
201 for cat_code, (cat_name, var_codes) in HIERARCHY.items():
202     local_w = group_weights[cat_code]
203     for vc in var_codes:
204         global_weights[vc] = Wc[cat_code] * local_w[vc]
205         variable_meta[vc] = cat_name
206
207 total_w = sum(global_weights.values()) # sanity check, should be ~1
208
209 # STEP 5: performance matrix
210 perf = read_performance_matrix(wb["Performance_Matrix"])
211 missing = [vc for vc in global_weights if vc not in perf]
212 if missing:
213     raise ValueError(f"Performance_Matrix is missing rows for: {missing}")
214
215 # STEP 6: scenario scores + ranking
216 scores = {scen: {} for scen in SCENARIOS}
217 for scen in SCENARIOS:
218     for arch in ARCHETYPES:
219         scores[scen][arch] = sum(global_weights[vc] * perf[vc][scen][arch] for vc in
global_weights)
220
221 rankings = {scen: sorted(ARCHETYPES, key=lambda a: scores[scen][a], reverse=True) for
scen in SCENARIOS}
222 top_picks = {scen: rankings[scen][0] for scen in SCENARIOS}
223 robust = len(set(top_picks.values())) == 1
224
225 # Category-level performance/contribution, for the dashboard charts.
226 # performance = local_w . P (0-1, comparable across categories);
227 # contribution = Wc * performance (sums to the total score across categories).
228 category_performance = {scen: {} for scen in SCENARIOS}
229 category_contribution = {scen: {} for scen in SCENARIOS}
230 for scen in SCENARIOS:
231     for cat_code, (cat_name, var_codes) in HIERARCHY.items():
232         local_w = group_weights[cat_code]
233         perf_by_arch = {arch: sum(local_w[vc] * perf[vc][scen][arch] for vc in
var_codes)
234                         for arch in ARCHETYPES}
235         category_performance[scen][cat_code] = perf_by_arch
236         category_contribution[scen][cat_code] = {arch: Wc[cat_code] * v for arch, v in
perf_by_arch.items()}
237
238 # Structural profile (Baseline only, not an AHP output): unweighted mean of
239 # the Performance Matrix scores per category, independent of expert judgement.
240 structural_profile = {
241     cat_code: {arch: sum(perf[vc]["Baseline"][arch] for vc in var_codes) /
len(var_codes)
242                for arch in ARCHETYPES}
243     for cat_code, (cat_name, var_codes) in HIERARCHY.items()
244 }
245
246 return {
247     "individual_report": individual_report,
248     "group_report": group_report,
249     "overall_cr_ok": overall_cr_ok,
250     "category_weights": Wc,
251     "local_weights": {k: v for k, v in group_weights.items() if k != "CAT"},
252     "global_weights": global_weights,
253     "variable_meta": variable_meta,
254     "total_global_weight_check": total_w,
255     "performance_matrix": perf,
256     "scores": scores,
257     "rankings": rankings,
258     "top_picks": top_picks,
259     "robust": robust,
260     "category_performance": category_performance,
261     "category_contribution": category_contribution,
262     "structural_profile": structural_profile,
263 }

```

```

264
265
266 def print_report(results):
267     print("=" * 70)
268     print("STEP 1 - INDIVIDUAL EXPERT CONSISTENCY (CR <= 0.10)")
269     print("=" * 70)
270     for r in results["individual_report"]:
271         print(f" {r['expert']:<10} matrix={r['matrix']:<4} n={r['n']} "
272               f"CR={r['CR']:.4f} -> {r['flag']}")
273
274     print()
275     print("=" * 70)
276     print("STEP 3 - GROUP (AGGREGATED) CONSISTENCY (CR <= 0.10)")
277     print("=" * 70)
278     for r in results["group_report"]:
279         print(f" matrix={r['matrix']:<4} n={r['n']} CR={r['CR']:.4f} -> {r['flag']}")
280     print(f"\n Overall group consistency: {'OK' - all matrices consistent' if
results['overall_cr_ok'] else 'REVIEW NEEDED'})")
281
282     print()
283     print("=" * 70)
284     print("STEP 4 - CATEGORY WEIGHTS (Wc)")
285     print("=" * 70)
286     for code, w in sorted(results["category_weights"].items(), key=lambda x: -x[1]):
287         name = HIERARCHY[code][0]
288         print(f" {code:<4} {name:<30} {w:.4f}")
289
290     print()
291     print(" GLOBAL VARIABLE WEIGHTS (Wi = Wc x wi|c), sorted descending")
292     for code, w in sorted(results["global_weights"].items(), key=lambda x: -x[1]):
293         print(f" {code:<6} ({results['variable_meta'][code]:<28}) {w:.4f}")
294     print(f"\n Sum check (should be ~1.0000): {results['total_global_weight_check']:.4f}")
295
296     print()
297     print("=" * 70)
298     print("STEP 6 - SCENARIO SCORES AND RANKING")
299     print("=" * 70)
300     for scen in SCENARIOS:
301         s = results["scores"][scen]
302         rank = results["rankings"][scen]
303         print(f" {scen:<10} Product={s['Product']:.4f} Use={s['Use']:.4f} "
304               f"Result={s['Result']:.4f} -> Ranking: {' > '.join(rank)}")
305
306     print()
307     print(f" Recommended archetype per scenario: {results['top_picks']}")
308     print(f" Recommendation is {'ROBUST (same archetype in all 3 scenarios)' if
results['robust'] else 'SENSITIVE (recommendation changes across scenarios)'}")
309
310
311 def export_report(results, output_path):
312     """Writes a readable Excel report of all pipeline outputs."""
313     from openpyxl.styles import Font, PatternFill, Alignment
314     wb = openpyxl.Workbook()
315     HEADER_FILL = PatternFill("solid", fgColor="305496")
316     HEADER_FONT = Font(bold=True, color="FFFFFF")
317     BOLD = Font(bold=True)
318
319     def style_header(ws, headers):
320         for c, h in enumerate(headers, start=1):
321             cell = ws.cell(row=1, column=c, value=h)
322             cell.font = HEADER_FONT
323             cell.fill = HEADER_FILL
324             cell.alignment = Alignment(horizontal="center")
325
326     # Individual CR
327     ws = wb.active
328     ws.title = "Individual_CR"
329     style_header(ws, ["Expert", "Matrix", "n", "Lambda_max", "CR", "Flag (threshold
0.10)"])
330     for r, row in enumerate(results["individual_report"], start=2):
331         ws.cell(row=r, column=1, value=row["expert"])

```

```

332     ws.cell(row=r, column=2, value=row["matrix"])
333     ws.cell(row=r, column=3, value=row["n"])
334     ws.cell(row=r, column=4, value=round(row["lambda_max"], 4))
335     ws.cell(row=r, column=5, value=round(row["CR"], 4))
336     ws.cell(row=r, column=6, value=row["flag"])
337
338     # Group CR
339     ws = wb.create_sheet("Group_CR")
340     style_header(ws, ["Matrix", "n", "Lambda_max", "CR", "Flag (threshold 0.10)"])
341     for r, row in enumerate(results["group_report"], start=2):
342         ws.cell(row=r, column=1, value=row["matrix"])
343         ws.cell(row=r, column=2, value=row["n"])
344         ws.cell(row=r, column=3, value=round(row["lambda_max"], 4))
345         ws.cell(row=r, column=4, value=round(row["CR"], 4))
346         ws.cell(row=r, column=5, value=row["flag"])
347     ws.cell(row=len(results["group_report"]) + 3, column=1,
348           value=f"Overall group consistency: {'OK' if results['overall_cr_ok'] else
'REVIEW NEEDED'}").font = BOLD
349
350     # Weights
351     ws = wb.create_sheet("Weights")
352     style_header(ws, ["Level", "Code", "Macro-category", "Weight"])
353     r = 2
354     for code, w in sorted(results["category_weights"].items(), key=lambda x: -x[1]):
355         ws.cell(row=r, column=1, value="Category")
356         ws.cell(row=r, column=2, value=code)
357         ws.cell(row=r, column=3, value=HIERARCHY[code][0])
358         ws.cell(row=r, column=4, value=round(w, 4))
359         r += 1
360     for code, w in sorted(results["global_weights"].items(), key=lambda x: -x[1]):
361         ws.cell(row=r, column=1, value="Variable (global Wi)")
362         ws.cell(row=r, column=2, value=code)
363         ws.cell(row=r, column=3, value=results["variable_meta"][code])
364         ws.cell(row=r, column=4, value=round(w, 4))
365         r += 1
366
367     # Scores
368     ws = wb.create_sheet("Scenario_Scores")
369     style_header(ws, ["Scenario", "Product", "Use", "Result", "Ranking", "Recommended"])
370     for r, scen in enumerate(SCENARIOS, start=2):
371         s = results["scores"][scen]
372         ws.cell(row=r, column=1, value=scen)
373         ws.cell(row=r, column=2, value=round(s["Product"], 4))
374         ws.cell(row=r, column=3, value=round(s["Use"], 4))
375         ws.cell(row=r, column=4, value=round(s["Result"], 4))
376         ws.cell(row=r, column=5, value=" > ".join(results["rankings"][scen]))
377         ws.cell(row=r, column=6, value=results["top_picks"][scen])
378     ws.cell(row=len(SCENARIOS) + 3, column=1,
379           value=f"Recommendation is {'ROBUST' if results['robust'] else 'SENSITIVE'} to
CRM price-volatility scenario.")
380
381     for sheet in wb.worksheets:
382         for col in sheet.columns:
383             length = max((len(str(c.value)) for c in col if c.value is not None),
default=10)
384             sheet.column_dimensions[col[0].column_letter].width = min(max(length + 2, 10),
45)
385
386     wb.save(output_path)
387
388
389 if __name__ == "__main__":
390     results = run("AHP_Data_Template.xlsx")
391     print_report(results)
392     export_report(results, "report.xlsx")

```

Listing B.1 - Source code of the AHP-based Decision Support System (ahp_analysis.py).

B.3 Program Output

Running the code on the case-study workbook produces two outputs. The first is the console report reproduced in Listing B.2, which lists the individual and group CRs, the category and global weights, and the scores and rankings of the three archetypes under each market condition. The second is the Excel workbook report.xlsx, whose four sheets are shown in Figs. B.7 to B.10: Individual_CR reports $\lambda_{\max}$, CR and consistency flag of each of the 28 individual matrices; Group_CR reports the same quantities for the seven aggregated matrices, together with the overall consistency check; Weights lists the category weights W_c and the global weights W_i of the eighteen decision variables in descending order; and Scenario_Scores reports the scores of the three archetypes, their ranking and the recommended archetype under each market condition, with the robustness check. The values coincide with those reported in Chapter 5, net of rounding. Since the input workbook contains the judgements after revision, the individual CRs of the revised matrices correspond to the second value reported for them in Tab. 5.1. Only cosmetic inconsistencies in variable and file names present in the code as originally executed have been corrected here for clarity; no computation or reported result is affected. The Organisational/Capabilities matrices ($n = 2$), for which the Consistency Ratio is zero by construction, are reported by the program as $CR = 0.000$ with flag OK; these are the same matrices marked n/a in Tab. 5.1, where they are excluded from the count of matrices subject to the 0.10 threshold.

```

=====
STEP 1 - INDIVIDUAL EXPERT CONSISTENCY (CR <= 0.10)
=====
Expert_1  matrix=CAT  n=6  CR=0.0878  -> OK
Expert_1  matrix=MKT  n=3  CR=0.0000  -> OK
Expert_1  matrix=ECO  n=4  CR=0.0395  -> OK
Expert_1  matrix=TEC  n=3  CR=0.0109  -> OK
Expert_1  matrix=ORG  n=2  CR=0.0000  -> OK
Expert_1  matrix=VCE  n=3  CR=0.0109  -> OK
Expert_1  matrix=REG  n=3  CR=0.0079  -> OK
Expert_2  matrix=CAT  n=6  CR=0.0724  -> OK
Expert_2  matrix=MKT  n=3  CR=0.0158  -> OK
Expert_2  matrix=ECO  n=4  CR=0.0225  -> OK
Expert_2  matrix=TEC  n=3  CR=0.0559  -> OK
Expert_2  matrix=ORG  n=2  CR=0.0000  -> OK
Expert_2  matrix=VCE  n=3  CR=0.0251  -> OK
Expert_2  matrix=REG  n=3  CR=0.0332  -> OK
Expert_3  matrix=CAT  n=6  CR=0.0656  -> OK
Expert_3  matrix=MKT  n=3  CR=0.0692  -> OK
Expert_3  matrix=ECO  n=4  CR=0.0006  -> OK
Expert_3  matrix=TEC  n=3  CR=0.0000  -> OK
Expert_3  matrix=ORG  n=2  CR=0.0000  -> OK
Expert_3  matrix=VCE  n=3  CR=0.0332  -> OK
Expert_3  matrix=REG  n=3  CR=0.0559  -> OK
Expert_4  matrix=CAT  n=6  CR=0.0676  -> OK
Expert_4  matrix=MKT  n=3  CR=0.0332  -> OK
Expert_4  matrix=ECO  n=4  CR=0.0367  -> OK
Expert_4  matrix=TEC  n=3  CR=0.0559  -> OK
Expert_4  matrix=ORG  n=2  CR=0.0000  -> OK
Expert_4  matrix=VCE  n=3  CR=0.0061  -> OK
Expert_4  matrix=REG  n=3  CR=0.0332  -> OK

=====
STEP 3 - GROUP (AGGREGATED) CONSISTENCY (CR <= 0.10)
=====
matrix=CAT  n=6  CR=0.0225  -> OK
matrix=MKT  n=3  CR=0.0064  -> OK
matrix=ECO  n=4  CR=0.0078  -> OK
matrix=TEC  n=3  CR=0.0007  -> OK
matrix=ORG  n=2  CR=0.0000  -> OK

```

```
matrix=VCE n=3 CR=0.0026 -> OK
matrix=REG n=3 CR=0.0066 -> OK
```

Overall group consistency: OK - all matrices consistent

STEP 4 - CATEGORY WEIGHTS (Wc)

```
ECO Economic/Business Case      0.4836
REG Regulatory/Institutional     0.1526
MKT Market                      0.1370
TEC Technical/Product           0.0855
ORG Organisational/Capabilities  0.0804
VCE Value Chain/Ecosystem       0.0610
```

GLOBAL VARIABLE WEIGHTS ($W_i = W_c \times w_i|c$), sorted descending

```
ECO2 (Economic/Business Case ) 0.1614
ECO1 (Economic/Business Case ) 0.1591
ECO3 (Economic/Business Case ) 0.0908
MKT1 (Market                  ) 0.0827
ECO4 (Economic/Business Case ) 0.0723
REG2 (Regulatory/Institutional ) 0.0676
REG3 (Regulatory/Institutional ) 0.0561
ORG1 (Organisational/Capabilities ) 0.0533
TEC2 (Technical/Product       ) 0.0526
VCE2 (Value Chain/Ecosystem   ) 0.0395
MKT2 (Market                  ) 0.0320
REG1 (Regulatory/Institutional ) 0.0289
ORG2 (Organisational/Capabilities ) 0.0271
MKT3 (Market                  ) 0.0223
TEC3 (Technical/Product       ) 0.0169
TEC1 (Technical/Product       ) 0.0160
VCE3 (Value Chain/Ecosystem   ) 0.0112
VCE1 (Value Chain/Ecosystem   ) 0.0103
```

Sum check (should be ~1.0000): 1.0000

STEP 6 - SCENARIO SCORES AND RANKING

```
Baseline Product=0.7808 Use=0.7039 Result=0.5891 -> Ranking: Product > Use > Result
Low       Product=0.8054 Use=0.7568 Result=0.6196 -> Ranking: Product > Use > Result
High      Product=0.6462 Use=0.6583 Result=0.6115 -> Ranking: Use > Product > Result
```

Recommended archetype per scenario: {'Baseline': 'Product', 'Low': 'Product', 'High': 'Use'}
Recommendation is SENSITIVE (recommendation changes across scenarios)

Listing B.2 - Console output of the DSS for the case-study data.

	A	B	C	D	E	F
1	Expert	Matrix	n	Lambda_max	CR	Flag (threshold 0.10)
2	Expert_1	CAT	6	6.5444	0.0878	OK
3	Expert_1	MKT	3	3	0	OK
4	Expert_1	ECO	4	4.1067	0.0395	OK
5	Expert_1	TEC	3	3.0126	0.0109	OK
6	Expert_1	ORG	2	2	0	OK
7	Expert_1	VCE	3	3.0126	0.0109	OK
8	Expert_1	REG	3	3.0092	0.0079	OK
9	Expert_2	CAT	6	6.449	0.0724	OK
10	Expert_2	MKT	3	3.0183	0.0158	OK
11	Expert_2	ECO	4	4.0606	0.0225	OK
12	Expert_2	TEC	3	3.0649	0.0559	OK
13	Expert_2	ORG	2	2	0	OK
14	Expert_2	VCE	3	3.0291	0.0251	OK
15	Expert_2	REG	3	3.0385	0.0332	OK
16	Expert_3	CAT	6	6.4064	0.0656	OK
17	Expert_3	MKT	3	3.0803	0.0692	OK
18	Expert_3	ECO	4	4.0017	0.0006	OK
19	Expert_3	TEC	3	3	0	OK
20	Expert_3	ORG	2	2	0	OK
21	Expert_3	VCE	3	3.0385	0.0332	OK
22	Expert_3	REG	3	3.0649	0.0559	OK
23	Expert_4	CAT	6	6.4189	0.0676	OK
24	Expert_4	MKT	3	3.0385	0.0332	OK
25	Expert_4	ECO	4	4.099	0.0367	OK
26	Expert_4	TEC	3	3.0649	0.0559	OK
27	Expert_4	ORG	2	2	0	OK
28	Expert_4	VCE	3	3.007	0.0061	OK
29	Expert_4	REG	3	3.0385	0.0332	OK

Fig. B.7 - Individual_CR sheet of the output workbook: individual consistency ratios of the 28 expert matrices.

	A	B	C	D	E
1	Matrix	n	Lambda_max	CR	Flag (threshold 0.10)
2	CAT	6	6.1397	0.0225	OK
3	MKT	3	3.0074	0.0064	OK
4	ECO	4	4.021	0.0078	OK
5	TEC	3	3.0008	0.0007	OK
6	ORG	2	2	0	OK
7	VCE	3	3.0031	0.0026	OK
8	REG	3	3.0077	0.0066	OK
9					
10	Overall group consistency: OK				

Fig. B.8 - Group_CR sheet of the output workbook: consistency ratios of the seven aggregated matrices.

	A	B	C	D
1	Level	Code	Macro-category	Weight
2	Category	ECO	Economic/Business Case	0.4836
3	Category	REG	Regulatory/Institutional	0.1526
4	Category	MKT	Market	0.137
5	Category	TEC	Technical/Product	0.0855
6	Category	ORG	Organisational/Capabilities	0.0804
7	Category	VCE	Value Chain/Ecosystem	0.061
8	Variable (global Wi)	ECO2	Economic/Business Case	0.1614
9	Variable (global Wi)	ECO1	Economic/Business Case	0.1591
10	Variable (global Wi)	ECO3	Economic/Business Case	0.0908
11	Variable (global Wi)	MKT1	Market	0.0827
12	Variable (global Wi)	ECO4	Economic/Business Case	0.0723
13	Variable (global Wi)	REG2	Regulatory/Institutional	0.0676
14	Variable (global Wi)	REG3	Regulatory/Institutional	0.0561
15	Variable (global Wi)	ORG1	Organisational/Capabilities	0.0533
16	Variable (global Wi)	TEC2	Technical/Product	0.0526
17	Variable (global Wi)	VCE2	Value Chain/Ecosystem	0.0395
18	Variable (global Wi)	MKT2	Market	0.032
19	Variable (global Wi)	REG1	Regulatory/Institutional	0.0289
20	Variable (global Wi)	ORG2	Organisational/Capabilities	0.0271
21	Variable (global Wi)	MKT3	Market	0.0223
22	Variable (global Wi)	TEC3	Technical/Product	0.0169
23	Variable (global Wi)	TEC1	Technical/Product	0.016
24	Variable (global Wi)	VCE3	Value Chain/Ecosystem	0.0112
25	Variable (global Wi)	VCE1	Value Chain/Ecosystem	0.0103

Fig. B.9 - Weights sheet of the output workbook: category weights and global weights of the decision variables.

	A	B	C	D	E	F	G
1	Scenario	Product	Use	Result	Ranking	Recommended	
2	Baseline	0.7808	0.7039	0.5891	Product > Use > Result	Product	
3	Low	0.8054	0.7568	0.6196	Product > Use > Result	Product	
4	High	0.6462	0.6583	0.6115	Use > Product > Result	Use	
5							
6	Recommendation is SENSITIVE to CRM price-volatility scenario.						

Fig. B.10 - Scenario_Scores sheet of the output workbook: scores, ranking and recommended archetype under each market condition.