Prospective Consequential Life Cycle Assessment for the Built Environment: A Framework with Explicit End-of-Life Modelling and Application to a Curtain Wall

by

Marley Kathleen Loft Dowling

A thesis

presented to the University of Waterloo

in fulfilment of the

thesis requirement for the degree of

Master of Applied Science

in

Civil Engineering

Waterloo, Ontario, Canada, 2026

© Marley Kathleen Loft Dowling 2026

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

The overall premise of this thesis was first introduced through an extended abstract published in the abstract book for the SETAC Europe 26th LCA Symposium, Making LCA Meaningful: Good Data, Better Models, Sustainable Decisions, held 21–23 October 2024 in Gothenburg, Sweden. The thesis significantly extends that preliminary work through a comprehensive literature review, full methodological development, case study application, and expanded analytical components. Chapter 2 of this thesis is a literature review completed entirely by myself. The review structure, synthesis of the relevant literature, and identification of the research gap were independently developed and written by me.

Chapters 3 and 4 of this thesis, excluding the uncertainty propagation components, have been incorporated into a journal paper intended for submission. This paper is co-authored by myself, Matt Roberts, Costa Kapsis, and Chris Bachmann. I led the development of the research design and methodological approach. The programming, data collection, analysis, visualization, and drafting of the manuscript were completed by me. Matt Roberts, Costa Kapsis, and Chris Bachmann provided supervision, technical input, manuscript feedback, and editorial revisions.

The uncertainty propagation and sensitivity analysis presented in this thesis are also being developed into a second journal paper with the same co-author group. The methodological development, implementation, analysis, visualization, and primary drafting associated with this work were completed by me, with guidance, feedback, and editing support from the co-authors.

Unless otherwise stated, the work presented in this thesis, including the implementation of methods, analysis, visualization, and writing, was completed by me.

Abstract

Life cycle assessment (LCA) is increasingly used to guide low-carbon decisions in the built environment, yet many building LCAs remain limited by two gaps: (i) weak alignment between the assessment year and the background systems that actually supply energy and materials over time, and (ii) simplified end-of-life (EoL) modelling that treats recycling and recovery pathways as static, mature, and well-characterized. These limitations are especially consequential for high-impact building materials and façade components, where EoL choices can shift impacts across regions, markets, and time horizons.

This thesis develops and demonstrates a prospective consequential LCA (pCLCA) framework for evaluating circular EoL strategies in buildings under future system changes, integrating (1) scenario-consistent background evolution using integrated assessment model (IAM) pathways coupled to life-cycle inventory databases, and (2) explicit, parameterized foreground models for both contemporary and emerging EoL technologies The framework operationalizes technology maturity through transparent scaling and data-quality rules, and incorporates a structured treatment of uncertainty via a structured taxonomy spanning foreground parameters, background evolution, and scenario choice. A reproducible implementation pipeline is presented using open-source tooling and a modular database architecture to enable consistent scenario generation, auditing, and sensitivity analysis.

The case study application, focused on aluminium façade EoL pathways, shows that route rankings and net outcomes are often governed more by Stage D substitution effects than by C3–C4 treatment burdens alone. In these credit-dominant regimes, negative net results should be interpreted as conditional displacement consequences shaped by displaced-product choices, market boundaries, marginal supplier representations, and equivalence rules, rather than as intrinsic "negative emissions" of the treatment routes themselves. The results further demonstrate that receiving-system sensitivity is route dependent. Uncertainty and sensitivity analysis reinforce these interpretations by showing that some patterns remain robust, while others, can approach or cross sign boundaries depending on background and receiving-system conditions.

Overall, the thesis contributes a transparent, decision-relevant, and extensible prospective consequential LCA method for evaluating circular EoL strategies in buildings under future system change, while emphasizing the need for explicit claims governance, decomposition-based reporting, and cautious interpretation of net-beneficial results.

Acknowledgements

I would like to sincerely thank my supervisors, Chris Bachmann and Konstantinos Kapsis, for their steady support, thoughtful feedback, and continued encouragement from the earliest stages of this research through to its completion; their mentorship has been foundational to both the quality of this thesis and my growth as a researcher.

I am also deeply grateful to Dr. Matt Roberts at the Center for the Built Environment, University of California, Berkeley, for his invaluable guidance and insight throughout this work—particularly in strengthening my approach to life cycle assessment, systems thinking, and modelling decisions.

This work was supported through funding from Mitacs Accelerate # IT32723 and Adaptis Technologies Inc.

Table of Contents

Author's Declaration ii
Statement of Contributions iii
Abstract iv
Acknowledgements v
List of Figures xi
List of Tables xiv
List of Abbreviations xvi
1 Introduction 1
1.1 Research Scope, Objectives and Limitations 2
2 Literature Review 5
2.1 Causality in Life Cycle Assessments 5
2.2 Temporal Considerations in Life Cycle Assessments 7
2.2.1 Prospective LCA 9
2.3 Technological Considerations in Life Cycle Assessments 9
2.3.1 Ex-Ante LCA 10
2.4 Characterizing Uncertainty & Variability in LCA 11
2.4.1 Model Uncertainty in LCA 13
2.4.2 Scenario Uncertainty in LCA 14
2.4.2.1 Sources of Scenario Uncertainty 15
2.4.2.2 Socioeconomic and Policy Pathways 20
2.4.2.3 Interdependencies and Structural Scenario Design 23
2.4.3 Parameter Uncertainty in LCA 25
2.4.3.1 Sources of Parameter Uncertainty 26
2.4.4 Variability in LCA 31
2.4.5 Interplay Between Variability and Uncertainty in LCA 31
2.4.6 Implications for Prospective LCA 32
2.5 Addressing Uncertainty & Variability in pLCA 32
2.5.1 Addressing Model Uncertainty in LCA 33
2.5.2 Addressing Scenario Uncertainty in LCA 35
2.5.2.1 Addressing Sources of Scenario Uncertainty 35

2.5.3 Addressing Parameter Uncertainty in LCA 44
2.5.3.1 Addressing Sources of Parameter Uncertainty 44
2.5.4 Addressing Variability in LCA 48
2.6 Summary & Research Gap 49
3 Methods 51
3.1 Methodological Framework Overview 51
3.2 Step 1: Establish Accounting Scope and Baseline Conventions 53
3.2.1 LCA Goal Definition & Objectives 53
3.2.2 LCA Scope & System Boundaries 54
3.2.2.1 Accounting Conventions 54
3.2.2.2 Admission Policy (applied to emerging routes only) 55
3.3 Step 2: Construct Scenario-Consistent Backgrounds 55
3.3.1 Scenario Selection and Bounding Uncertainty 56
3.3.2 Prospective Pathway Integration 56
3.3.3 Data Processing & Representation of Scenario Drivers via Premise 57
3.3.4 Limitations of Scenario Translation 58
3.4 Step 3: Define Market Boundaries and Compute Marginal Supplier Mixes 58
3.4.1 Step 3-1: Identifying the Scale and Time Horizon of the Change 59
3.4.2 Step 3-2: Identifying the Limits of a Market 59
3.4.3 Step 3-3: Identifying Trends in the Volume of a Market 60
3.4.4 Step 3-4: Identifying Changes in Supply & Demand 61
3.5 Step 4: Model Foreground EoL Pathways and Stage D Substitution 62
3.5.1 Foreground EoL Route Modelling 62
3.5.2 Proxy Foreground Processes 63
3.5.2.1 Utility Regionalization of Foreground Proxies 63
3.5.3 Evidence-Based Scaling 63
3.5.4 Stage D Substitution 64
3.5.4.1 Equivalence & Conditioning 65
3.5.5 Deterministic LCA 66
3.5.5.1 Contribution Analysis & Hotspot Features 66
3.6 Step 5: Propagate Uncertainty and Test Sensitivity 66

3.6.1 Uncertainty Taxonomy 67
3.6.2 Foreground Parameter Uncertainty 67
3.6.3 Evidence-Based Range Elicitation 67
3.6.3.1 Edge Cases and Sparse-Information Rules 68
3.6.4 Monte Carlo Conditioning 69
3.6.5 Hotspot Features via Sensitivity Analysis 69
3.6.6 Scenario Uncertainty and Reporting 70
3.7 Summary 70
4 Case Study: Application to a Curtain Wall 72
4.1 Overview & Reader's Map 72
4.2 Step 1: Accounting Scope and Baseline Conventions 73
4.2.1 Functional Unit 73
4.2.2 Scope & System Boundaries 74
4.2.2.1 Geographic 74
4.2.2.2 Temporal 75
4.2.3 Bill of Materials 75
4.2.4 Admitted Route Register and Emerging-Technology Selection 76
4.2.4.1 Admitted Aluminium End-of-Life Routes 77
4.2.5 Contemporary Snapshot 77
4.3 Step 2: Scenario-Consistent Backgrounds 78
4.3.1 Scenario Set & Database Instantiation 78
4.3.2 Background Construction & Freeze (v1.0 background set) 79
4.3.3 Post-Build Quality Assurance (QA) and "No-Edit" Outcome 79
4.3.3.1 Integrity and Linkage QA (background self-consistency) 79
4.3.3.2 Project-Wide Orphan Exchange Prevention 80
4.3.3.3 Geographic Anchoring: Bridges Not Required 80
4.3.3.4 Versioning and Freezing Record 80
4.3.4 Limitations 80
4.4 Step 3: Defining Market Boundaries & Computing Marginal Supplier Mixes 81
4.4.1 Step 3-1: Scale and Time Horizon of Change 82
4.4.2 Step 3-2: Limits of a Market 83

4.4.2.1 Reuse 83
4.4.2.2 Conventional Recycling 84
4.4.2.3 Mid-TRL Emergent Recycling 84
4.4.2.4 Low-TRL Emergent Recycling 84
4.4.3 Step 3-3: Trends in the Volume of a Market 85
4.4.3.1 Market Trends for Contemporary Aluminium Billets 86
4.4.3.2 Market Trends for Contemporary Extruded-Aluminium 89
4.4.3.3 Market Trends for Contemporary Industrial Gaseous Hydrogen 89
4.4.4 Step 3-4: Changes in Supply and Demand 90
4.4.4.1 Market Trends for Contemporary Aluminium Billets 93
4.4.4.2 Market Trends for Contemporary Extruded-Aluminium 94
4.4.4.3 Market Trends for Contemporary Industrial Gaseous Hydrogen 94
4.4.4.4 Contemporary Electricity Supply 95
4.5 Step 4: Model Foreground EoL Pathways and Stage D Substitution 96
4.5.1 Foreground EoL Route Modelling 96
4.5.1.1 Route Mechanics 97
4.5.2 Proxy Foreground Processes 99
4.5.3 Evidence-Based Scaling 99
4.5.3.1 Multi-Step Friction-Stir Consolidation (MS-FSC): Evidence-Based Scaling 100
4.5.3.2 Aluminium Hydrolysis: Evidence-Based Scaling 101
4.5.4 Stage D Substitution 102
4.5.4.1 Product-to-Market Mapping 103
4.5.4.2 Equivalence & Conditioning 103
4.5.5 Deterministic Results 105
4.5.5.1 Stage D as a Qualified Substitution Claim 107
4.5.5.2 Interpreting Dominance 109
4.5.5.3 Landfilling as a Baseline 110
4.5.5.4 Within-Sector Displacement: Reuse and Conventional Recycling 111
4.5.5.5 Cross-Sector Displacement: Hydrolysis and Receiving-System Dependence 112
4.5.5.6 Emerging Technology: MS-FSC and Compounded Dependence 114

4.5.5.7 Cross-Indicator Screening: Burden Shifting Under a Stage D-Governed Signal 116
4.5.5.8 Deterministic Stress-Test Suite Derived from Dominance Screening 118
4.5.5.9 What Deterministic Results Establish 119
4.5.5.10 Implications for Prospective Module D Reporting and Interpretation 119
4.6 Step 5: Uncertainty Propagation and Sensitivity Evidence 120
4.6.1 Executed Uncertainty Layers and Interpretation Scope 120
4.6.2 Robustness Screen: Sign Stability and Dominance Structure 121
4.6.3 Convergence and Adequacy 123
4.6.4 Sensitivity Hotspots 125
4.6.5 Implications for Interpreting Deterministic Results 127
5 Conclusions 128
5.1 Methodological Contributions 128
5.2 Key Findings 129
5.3 Scope and Limitations 130
5.4 Future Work 131
References 132
Appendix A — Stage D Substitution Governance Pack 147
Appendix B — Foreground Models, Proxies, Parameters, and Regionalization 153
Appendix C — Background Database Construction, QA, and Reproducibility 173
Appendix D — Markets, Marginal Suppliers, and Data Provenance 181
Appendix E — Additional LCIA Categories and Qualitative Stability 190
Appendix F — Uncertainty Propagation and Sensitivity Evidence 199
Appendix G — Code Repository 226

List of Figures

Figure 2-1: An Illustrative Time-Maturity Map of Life Cycle Assessment with Causality as an Orthogonal Dimension (derived from Arvidsson et al. (2023)) 5
Figure 2-2: Classification of uncertainty and variability in life cycle assessment (LCA). The diagram separates uncertainty arising from imperfect knowledge or modelling choices—grouped as parameter, scenario, and model uncertainty—from variability, which reflects real-world heterogeneity across space, time, and sources/objects..... 12
Figure 2-3: Conceptual schematic distinguishing three LCA time horizons and their dominant uncertainty pathways. (A) The assessment horizon governs projection depth of the receiving system; longer horizons increase dependence on scenario assumptions, illustrated by divergence across plausible receiving-system pathways and an expanding outcome space (schematic, not LCIA results). (B) The inventory horizon defines the time window over which life-cycle flows are tracked; temporally concentrated phases (e.g., construction and end-of-life processing) may appear as pulses relative to the full lifetime, and delayed flows (e.g., slow-release emissions or extended circular loops) can extend beyond the cutoff, creating truncation risk when omitted. (C) The characterization horizon is a metric window length (e.g., GWP20/100/500) applied after emissions occur, influencing temporal weighting and metric sensitivity. Curves and thresholds are illustrative and not to scale..... 17
Figure 2-4: The sequence of information used to project future levels of climate change (derived from Riahi et al., 2017)..... 22
Figure 2-5: Illustrative graphic of the relationship between technology diffusion, knowledge accumulation, and design degrees of freedom across time in prospective LCA, with an indicative trajectory of parameter uncertainty due to scaling/extrapolation. TRL/MRL is shown as a readiness concept relevant at the assessment time $t0$, while modelled performance at tf typically assumes manufacturing readiness (adapted from Arvidsson et al., 2018)..... 30
Figure 2-6: Schematic of the IAM–LCA coupling workflow used to generate prospective background databases derived from (Sacchi et al., 2022). (1) Integrated assessment model outputs (REMIND, IMAGE, GCAM, MESSAGEix, TIAM-UCL) provide scenario trajectories that (2) are applied to the baseline ecoinvent database via premise to create policy- and scenario-specific inventories (e.g., no climate policy, Paris Agreement target, nationally determined contributions, biofuel push), optionally augmented with additional inventories for emerging technologies. (3) The resulting scenario-specific databases are then used in LCA tools (e.g., Brightway/Activity Browser, SimaPro, openLCA, Python)..... 40
Figure 2-7: Four-step heuristic for identifying long-term marginal supplier responses in consequential LCA (derived from (B. Weidema, 2010))..... 41
Figure 3-1: Proposed workflow for prospective consequential LCA (pCLCA) of end-of-life pathways for the built environment: baseline accounting conventions, IAM-derived background scenarios, marginal-mix computation, deterministic Stage D substitution results, and uncertainty/sensitivity analysis..... 53

Figure 3-2: Whole-life carbon assessment scope and module structure. Life cycle are shown following the building life cycle module convention (A1–A3 product stage, A4–A5 construction, B use stage, C end-of-life, and Module D beyond the system boundary), with included and excluded modules delineated by the system boundary. 54
Figure 4-1: High level schematic curtain wall panel with visible major assemblies labelled 74
Figure 4-2: Conceptual illustration of the capital replacement rate (CRR) used to distinguish suppliers that are net expanding (green dashed trajectories; “modern” technologies contributing to marginal supply) from those that are net contracting (red dashed trajectories; “old” technologies). The black line represents the CRR threshold (e.g., 3.33% per year for a 30-year lifetime), showing the minimum growth needed to exceed replacement and indicate effective capacity expansion over time. Derived from (Consequential-LCA, 2020). 92
Figure 4-3: Illustrative Global Aluminium Flow (from Boin & Bertram, 2005). All units measured in million tonnes per year. Values might not add up due to rounding. a–Aluminum in dross; b–not taken into account in statistics; c–such as powder, paste, and deoxidation aluminum (metal property is lost); d–area of current research to identify final aluminum destination (reuse, recycling, or landfilling);e–includes, depending on the ore, between 30% and 50% alumina; f–includes, on a global average, 52% aluminum. 98
Figure 4-4: Illustrative Single & Multi-Step Friction Stir Consolidation Process (taken from Ingarao et al., 2023) 99
Figure 4-5: Illustrative Aluminium Hydrolysis Process (derived from Urbonavicius et al., 2024) 99
Figure 4-6: Multi-Step Friction Stir Consolidation (MS-FSC) Contemporary System Boundaries & Crediting (Unit Process) 104
Figure 4-7: Aluminium Hydrolysis Contemporary System Boundaries & Crediting (Unit Process) 105
Figure 4-8: Decomposed deterministic climate change results (GWP100, kg CO_2-eq): C3–C4, Stage D, and net, across EoL routes for 2025 contemporary and 2050 prospective SSP variants 106
Figure 4-9: Dominance Ratio vs Net Climate Change Interpretation Map 110
Figure 4-10: Decomposed deterministic climate change results (kg CO_2-eq): C3–C4 only, across EoL routes for 2025 contemporary and 2050 prospective SSP variants.....111
Figure 4-11:Heat map of Sign Consistency vs GWP for Non-Climate Midpoints 117
Figure 4-12: Contemporary background exchange uncertainty: net robustness across routes (GWP100) 123
Figure 4-13: Prospective foreground exchange uncertainty: net robustness across routes and scenarios (GWP100) 123
Figure 4-14: Hydrolysis driver plot showing the relationship between NaOH recirculated makeup fraction (f_makeup) and net outcomes across SSP backgrounds (Monte Carlo draws shown with SSP overlays)..... 126

Figure 4-15: MS-FSC driver plot showing the relationship between credit realization share (pass_share) and net outcomes across SSP backgrounds..... 127
Figure F-1: Contemporary background exchange uncertainty: net robustness across routes (GWP100)..... 200
Figure F-2: Prospective foreground-only uncertainty: net GWP100 distributions across routes and SSP pathways (p5–p95). 203
Figure F-3: Hydrolysis key drivers (SSP1-VLLO, 2050; gate-basis net GWP; decile medians with 10–90 bands)..... 216
Figure F-4: Hydrolysis key drivers (SSP2-M, 2050; gate-basis net GWP; decile medians with 10–90 bands)..... 216
Figure F-5: Hydrolysis key drivers (SSP5-H, 2050; gate-basis net GWP; decile medians with 10–90 bands)..... 217
Figure F-6: MS-FSC key drivers (SSP1-VLLO_2050; gate-basis net GWP; decile medians with 10–90 bands)..... 218
Figure F-7: MS-FSC key drivers (SSP2-M_2050; gate-basis net GWP; decile medians with 10–90 bands)..... 218
Figure F-8: MS-FSC key drivers (SSP5-H_2050; gate-basis net GWP; decile medians with 10–90 bands)..... 218
Figure F-9: Hydrolysis SSP overlay: net vs hydrogen purification recovery (R_PSA) (gate-basis net GWP) 220
Figure F-10: Hydrolysis SSP overlay: net vs scrap preparation yield (Y_prep) (gate-basis net GWP) 220
Figure F-11: Hydrolysis SSP overlay: net vs reacted aluminium fraction (X_Al) (gate-basis net GWP) 221
Figure F-12: Hydrolysis SSP overlay: net vs caustic make-up fraction (f_makeup) (gate-basis net GWP) 221
Figure F-13: Hydrolysis SSP overlay: net vs metallic aluminium fraction in prepared feed (f_Al) (gate-basis net GWP)..... 222
Figure F-14: MS-FSC SSP overlay: net vs billet qualification/eligibility share (pass_share) (gate-basis net GWP)..... 223
Figure F-15: MS-FSC SSP overlay: net vs transition overhead retention fraction (f_transition) (gate-basis net GWP)..... 223

List of Tables

Table 2-1: Comparison of Prospective LCA (pLCA) and Ex-Ante LCA 11
Table 4-1: Curtain wall bill of materials aligned to major assemblies shown in Figure 10. Components are organized by assembly/zone (vision glazing, seals/gaskets, frame and thermal break, spandrel/backpan zone, flashings/trim, and subframing/backing) with material designation and quantity reported per square metre of installed curtain wall (kg/m²). 75
Table 4-2: Primary aluminium (billet-equivalent proxy) – window volumes and expansion signal by historic domestic producer (2018-2023) 87
Table 4-3: Contemporary marginal supplier mix for the aluminium billet market (Canada; 2018–2023 expansion-signal basis) 93
Table 4-4: Canada-wide marginal electricity mix derived from 2025–2030 projected generation changes (technology-level shares) 96
Table 4-5: Prospective (2050) Central Scaling Parameters for Aluminium Hydrolysis (EoL Treatment) 101
Table 4-6: Decomposed deterministic climate change results (GWP100, kg CO_2-eq): C3–C4, Stage D, and net, across EoL routes for 2025 contemporary and 2050 prospective SSP variants 107
Table 4-7: Hydrolysis receiving-system leverage and break-even credit scaling (derived from deterministic C3–C4, Stage D, displaced-hydrogen intensity, and unit electricity burdens used in each scenario-year). 113
Table 4-8: MS-FSC deterministic split and receiving-system context (primary method, GWP100). 115
Table 4-9: Robustness Screen for Uncertainty Propagation (GWP100; medians and p5–p95; dominance at p50) 122
Table 4-10: Convergence evidence (full diagnostics in Appendix F): 125
Table A-1: Stage-D substitution-claim ledger for retained routes 149
Table A-2: Product-to-market mapping for creditable outputs 150
Table A-3: Governance-Variant Register 151
Table B-1: Targeted-review longlist and screening register for candidate emerging aluminium EoL technologies 154
Table B-2: Full bill of materials for the representative curtainwall panel (~33m^2 surface area) 156
Table B-3: Foreground route wrapper summary 159
Table B-4: Proxy registry 161
Table B-5: Principal MS-FSC baseline parameters 163
Table B-6: Principal hydrolysis baseline parameters 164
Table B-7: Utility and provider-alignment summary 166
Table B-8: Non-priority material treatment 170
Table B-9: Selected non-priority material net GWP100 results (kg CO_2e/m²) 170
Table C-1: Contemporary background snapshot record 174
Table C-2: Contemporary archived artifacts and SHA-256 hashes 175

Table C-3: Prospective scenario background structural QA summary 177
Table C-4: Prospective scenario background identifiers and integrity records..... 178
Table C-5: Quality-assurance checks applied 179
Table D-1: Market boundary register for retained displaced products and services 182
Table D-2: Evidence families used to define markets and marginal suppliers..... 184
Table D-3: Contemporary marginal supplier mix for the aluminium billet market (Canada; 2018–2023 basis) 185
Table D-4: Canada-wide marginal electricity mix derived from 2025–2030 projected generation changes..... 187
Table E-1: ReCiPe 2016 Midpoint (H) indicators included in the screening set 191
Table E-2: Sign consistency versus GWP100 across the 17 non-climate midpoint indicators.. 193
Table E-3: Non-climate midpoint indicators with at least one sign reversal relative to GWP100 194
Table E-4: Rank-inversion summary relative to GWP100 195
Table F-1: Contemporary (2025) background exchange uncertainty (GWP100, kg CO_2-eq per FU). 202
Table F-2: Hydrolysis (2050) foreground-only uncertainty by SSP pathway (GWP100, kg CO_2-eq per FU). Iteration counts (n) reflect convergence-stopped Monte Carlo budgets; convergence diagnostics are reported in §4.7.4. 206
Table F-3: MS-FSC (2050) foreground-only uncertainty by SSP pathway (GWP100, kg CO_2-eq per FU). Iteration counts (n) reflect convergence-stopped Monte Carlo budgets; convergence diagnostics are reported in §4.7.4. 206
Table F-4: Foreground Uncertainty Inputs of Prospective Hydrolysis 207
Table F-5: Foreground Uncertainty Inputs of Prospective MS-FSC 208
Table F-6: Prospective FG-only convergence summary (qRMSE_rel criterion; primary method; quantile grid p5–p95; tol = 1%; stable hits ≥ 3)..... 212

List of Abbreviations

Abbreviation	*Meaning*
A1–A3	Product stage
A4–A5	Construction process stage
AB	Alberta
AEC	Architecture, Engineering, and Construction
Al	Aluminum
ALCA	Attributional Life Cycle Assessment
APOS	Allocation at the Point of Substitution
B1–B7	Use stage
BC	British Columbia
C1–C4	End-of-life stage
CA	Canada
CaGBC	Canada Green Building Council
CE	Circular Economy
CI	Confidence Interval
CLCA	Consequential Life Cycle Assessment
CO_2	Carbon dioxide
CO_2e	Carbon dioxide equivalent
D	Benefits and loads beyond the system boundary
DfD	Design for Disassembly
DQI	Data Quality Indicator
EAF	Electric Arc Furnace
EN	European Standard (Norme Européenne)

Abbreviation	*Meaning*
EoL	End-of-Life
EPD	Environmental Product Declaration
FU	Functional Unit
GHG	Greenhouse Gas
GWP	Global Warming Potential
GWP100	Global Warming Potential over a 100-year time horizon
H	Hierarchist perspective
IAM	Integrated Assessment Model
ISO	International Organization for Standardization
kg	Kilogram
kWh	Kilowatt-hour
LCA	Life Cycle Assessment
LCI	Life Cycle Inventory
LCIA	Life Cycle Impact Assessment
LHS	Latin Hypercube Sampling
MC	Monte Carlo
MFA	Material Flow Analysis
MS-FSC	Multi-Step Friction Stir Consolidation
ON	Ontario
pCLCA	Prospective Consequential Life Cycle Assessment
pLCA	Prospective Life Cycle Assessment
PSA	Pressure Swing Adsorption

Abbreviation	*Meaning*
QC	Quebec
RICS	Royal Institution of Chartered Surveyors
SSP	Shared Socioeconomic Pathway
SSP1-VLLO	Shared Socioeconomic Pathway 1, very-low/low-overshoot scenario
SSP2-M	Shared Socioeconomic Pathway 2, middle-of-the-road scenario
SSP5-H	Shared Socioeconomic Pathway 5, high-emissions scenario
tCO_2e	Tonnes of carbon dioxide equivalent
TRL	Technology Readiness Level

Chapter 1

1 Introduction

In the pursuit of sustainable development, the built environment stands as a defining sector, driving global carbon emissions and resource consumption. As urbanization accelerates, the demand for construction materials and energy intensifies, compounding environmental pressures. Addressing these challenges requires strategies that reduce environmental impacts, minimize waste, and promote resource efficiency—while also accounting for the evolving economic, societal, and technological landscape.

Life Cycle Assessments (LCAs) have become the leading method for evaluating the environmental impacts of construction practices, identifying hotspots, and revealing opportunities for improvement across the full life cycle of building materials and systems. However, conventional LCA approaches often fail to reflect the reality of a changing world, relying on static assumptions about the lifespan of construction assets and the environmental conditions that will persist over an entire assessment period. While this simplification makes implementation easier, it can introduce significant uncertainty, particularly in areas where technological progress is expected to drastically shift environmental outcomes.

In this regard, decisions should ideally be informed by an understanding of their environmental consequences, necessitating a more forward-looking approach. Consequential Life Cycle Assessment (CLCA) provides a valuable framework by quantitatively assessing the broader environmental implication of decisions rather than merely attributing impacts to existing systems. However, because these consequences unfold over time, it is critical to account for the expected evolution of the socio-economic context, which in turn influences the background system of an LCA. The background system in this context refers to the broader supply chains and supporting processes, such as energy generation, raw material extraction and waste management, that lie outside the direct control of the producer. Without incorporating these dynamic factors, long-term environmental predictions risk becoming detached from the reality of future conditions. This underscores the need for a prospective CLCA (pCLCA) approach, one that aligns foreground system assumptions (i.e., the processes directly controlled or influenced by the decision-maker) with the anticipated transformations of the background system to produce assessments that are not only accurate but also actionable in shaping sustainable decision-making.

Despite growing recognition of the need for temporal consistency in LCA, few studies have successfully integrated prospective CLCA with modified life cycle inventories (LCIs) that evolve over time (Maes et al., 2023, 2024). While some research has explored future scenario-based modifications of LCIs, many of these studies focus on general industrial processes rather than the built environment (Jones et al., 2017; Maes et al., 2021). Even among the studies that adopt a consequential modelling approach, the majority emphasize primary material production rather than end-of-life dynamics, leaving a critical gap in understanding how future waste management practices, recycling efficiencies, and secondary material markets will shape environmental outcomes. Given the long lifespan of buildings (i.e., 50 to 100+ years) and the significant role that construction and demolition (C&D) waste plays in global material flows, the lack of prospective CLCA studies in the built environment represents a major limitation in current research and practice.

1.1 Research Scope, Objectives and Limitations

This study distinguishes itself by applying prospective CLCA (pCLCA) within the built environment while also explicitly addressing end-of-life (EoL) material flows, a critical but underexplored area in LCA research. More specifically, the central contribution of this research is the methodological framework developed to align prospective background change, consequential market logic and end-of-life (EoL) pathway modelling, and uncertainty analysis within a single decision-relevant workflow. A case study is additionally executed to demonstrate the implementation and interpretive value of this framework. The resulting objectives of this study are thus to:

- Evaluate the integration of IAM-driven background system modifications into prospective CLCA at a specific point in time.
- Evaluate the role of emerging recycling technologies in shaping future end-of-life impacts for construction materials.
- Assess the influence of supplier competitiveness and marginal production shifts on LCA results, ensuring consistency between short- and long-term environmental forecasts.
- Explore the implications of scenario and parameter uncertainty in prospective modelling by identifying and characterizing the most consequential sources in both foreground and background systems, including technology performance, adoption rates, time horizons, and systemic transitions.

To meet these objectives, this study:

- Modifies the Ecoinvent consequential database to align with Canadian industry conditions, ensuring accurate representation of regional material supply chains;
- Integrates the IMAGE Integrated Assessment Model (IAM) projections to capture macroeconomic shifts, energy system transitions, and changes in material circularity; and
- Propagates uncertainty through probabilistic modelling approaches, using Monte Carlo simulations to account for both scenario-related variation (e.g., time horizon, background system evolution, supplier shifts) and parameter uncertainty (e.g., process efficiencies, emission factors, energy use), and evaluate their combined impact on LCA outcomes, with the executed design limited to contemporary background-exchange uncertainty and prospective foreground-parameter uncertainty within fixed SSP backgrounds.

By incorporating dynamic background change and structured treatment of scenario and parameter uncertainty, this study strengthens the temporal credibility of prospective CLCA in the built environment. The framework links end-of-life modelling, marginal supplier responses, and regionalized supply chains to produce decision-relevant results. These objectives are pursued through the development and demonstration of the framework on a curtain wall system. The choice prioritizes a complex, multi-material assembly where interfaces, fasteners, and adhesives constrain disassembly; service lives are uneven; and high-value materials such as aluminium and glass participate in evolving secondary markets—conditions that make it an efficient test bed for prospective CLCA.

Despite its contributions, this study is subject to several limitations within the scoping of a curtain wall system. A key limitation of this study is its regional specificity to the Canadian context, which constrains the generalizability of its findings to other geographic settings without comparable modelling adaptations. To improve relevance, the life cycle inventory was tailored to reflect

Canadian energy systems, material supply chains, and end-of-life pathways. However, these regional adjustments are layered onto a global background database, which was not originally designed to represent Canada's industrial or environmental conditions. This partial regionalization creates a potential mismatch between modelled and actual supply chains, particularly in sectors where localized data are sparse or aggregated. While these modifications improve accuracy relative to unadjusted global models, they also introduce residual uncertainty regarding the representativeness of background system dynamics.

Furthermore, while the study emphasizes end-of-life material flows and background system evolution, it does not explicitly model changes in construction techniques, building design innovations, or on-site practices, all of which can substantially influence life cycle impacts. Social, economic, and policy feedback, such as public acceptance of emerging recycling technologies, market penetration constraints, and labour or job creation implications, are not accounted for in the modelling. Although the Integrated Assessment Model (IAM) incorporates broad policy scenarios (e.g., climate targets, carbon pricing), this study does not simulate dynamic or region-specific policy responses that may influence technology adoption or circularity strategies over time.

The use of the IAM further introduces limitations as a direct result of its sectoral aggregation and coarse spatial resolution, which may overlook fine-grained variations in localized supply chains and material flows. Finally, the uncertainty analysis, while comprehensive in its treatment of scenario and parameter variability, does not extend normative or structural uncertainties embedded in prospective assumptions or system boundary choices which could further affect the interpretation of outcomes.

This thesis is organized into five chapters. Chapter 1 introduces the motivation, research scope, and objectives, highlighting the limitations of conventional life cycle assessment (LCA) in capturing long-term system dynamics and setting the stage for a prospective, consequential approach. Chapter 2 builds on this foundation by reviewing the literature on prospective and consequential life cycle assessment, with particular attention to methodological challenges such as uncertainty propagation, foreground and background system alignment, and the treatment of emerging technologies. Insights from the review directly informs Chapter 3, which presents the proposed methodological framework. This framework details the selection and customization of LCA scenarios, system boundaries, and modelling logic detailing end-of-life processes and power grid projections. Chapter 4 details the application of Chapter 3's developments to a case study centred on a curtain wall as an application demonstration of the proposed framework and documents the results of the pCLCA across modelled scenarios in addition to offering a critical discussion on the key findings, emphasizing trade-offs, modelling implications, and system expansion logic. Finally, Chapter 5 concludes the thesis by synthesizing insights from the prospective LCA analysis, highlighting the overall contributions, limitations, and implications for future research and practice in sustainable construction.

The appendices provide the detailed technical, evidentiary, and audit material that supports the main thesis chapters. Appendix A documents the Stage D substitution governance used in the study, Appendix B records the foreground modelling layer, Appendix C compiles supplementary case-study implementation records, including background freeze artifacts, manifests, and related reproducibility material. Appendix D documents the market-boundary definitions, marginal supplier calculations, trend evidence, and supporting worked. Appendix E provides supplementary

deterministic and cross-indicator result material used to assess burden shifting and indicator consistency beyond the primary climate-change discussion. Appendix F reports the uncertainty propagation and sensitivity-analysis framework in full, including parameter ranges, convergence diagnostics, and additional robustness outputs. Lastly, Appendix G provides the GitHub code repository link and accompanying implementation archive that support the models, scripts, processing steps, and reproducibility materials referenced throughout the preceding appendices.

Chapter 2

2 *Literature Review*

Recent work in life cycle assessment (LCA) has emphasized the importance of three interrelated dimensions: causality, technological maturity, and time (Arvidsson et al., 2023). The preeminent portion of this review draws on that framing to explore how LCA has developed to support more future-oriented assessments, focusing in particular on consequential and prospective approaches. These three dimensions and their interrelated nature are visually depicted in Figure 2-1.

This review also examines how these extensions give rise to new forms of uncertainty, both structural and parametric, especially when modelling long-term changes. Foundational contributions by Lloyd and Ries (2007) and Huijbregts (1998a, 1998b) are used to guide the second part of the review, which addresses how uncertainty and variability are introduced across modelling choices. This framing underpins the methodological direction taken in this thesis.

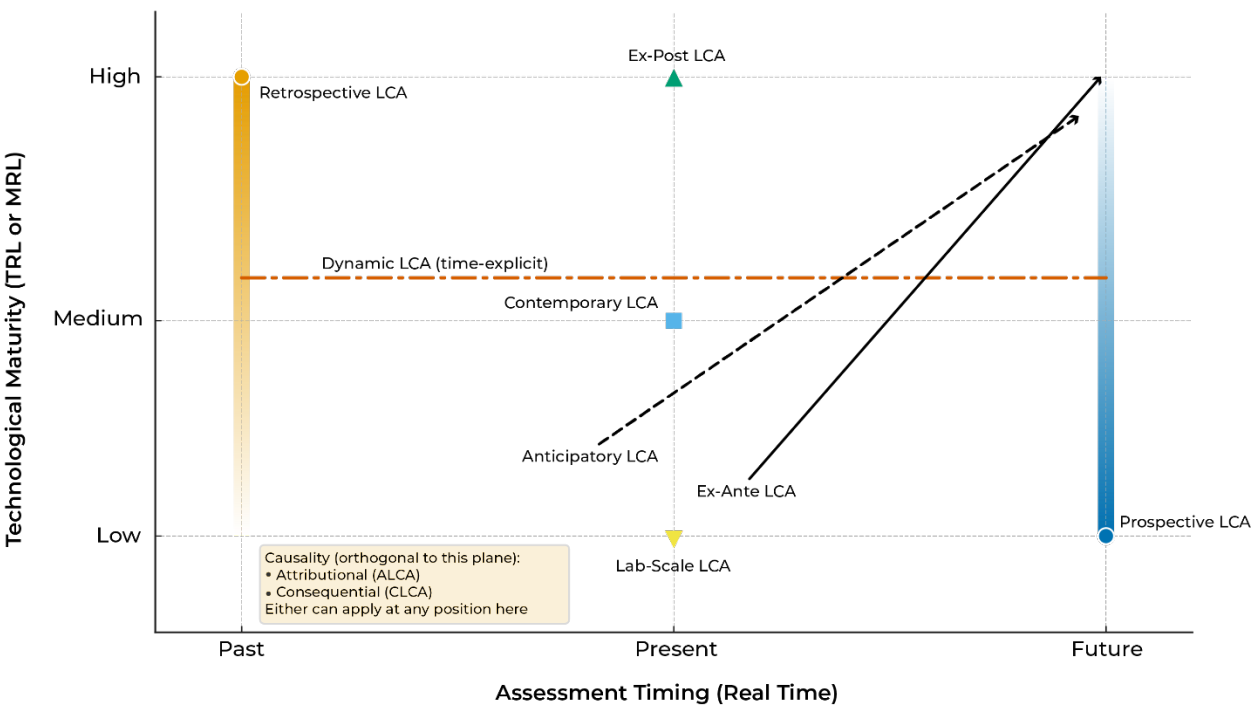


Figure 2-1*: An Illustrative Time-Maturity Map of Life Cycle Assessment with Causality as an Orthogonal Dimension (derived from Arvidsson et al. (2023))*

2.1 *Causality in Life Cycle Assessments*

Life Cycle Assessment (LCA) methodologies are commonly classified into Attributional LCA (ALCA) and Consequential LCA (CLCA), the distinction rooted in how they define system boundaries, model causal relationships, and allocate environmental burdens. ALCA quantifies the environmental impacts of a product or system within a fixed system boundary, typically using average data to represent existing market conditions (Ekvall, 2020). This approach is widely used for corporate sustainability reporting, product benchmarking, and regulatory

compliance, as it provides a consistent framework for comparing environmental performance across different products and industries (Schaubroeck, 2023).

ALCA has been criticized on two fronts. First, some scholars argue that ALCA is frequently misapplied to scenarios where CLCA would be more appropriate. ALCA is designed to describe the environmentally relevant physical flows associated with a product or system, whereas CLCA is intended to model how these flows change in response to a decision, capturing broader market and system-level consequences (Finnveden et al., 2009). Weidema et al. (2018) argue that the fixed system boundaries of ALCA inherently exclude indirect effects beyond the assessed system, meaning that any changes to a product or process can have consequences outside of what is accounted for. Secondly, concerns have been raised regarding ALCA's allocation methods, particularly the often arbitrary or artificial division of environmental burdens among co-products in multi-output processes (Schaubroeck, 2023). The authors highlight that the allocation of environmental burdens, whether through partitioning impacts across co-products or averaging data within a system, can mask the influence of broader supply chain shifts, policy interventions, or market dynamics. These allocation choices can significantly influence results and may not always reflect real-world economic or environmental relationships.

CLCA, however, is designed to evaluate the environmental implications of decisions by incorporating market-mediated effects and system-wide consequences (Brandão et al., 2024). Rather than assigning a share of global emissions to a product system, CLCA models how changes in demand or supply propagate through markets, influencing marginal production technologies and resource allocation (Mendoza Beltran et al., 2020). This broader perspective makes CLCA particularly valuable for assessing policy interventions, large-scale technological transitions, and shifts in consumer behaviour, where ripple effects extend beyond a single product or process.

Furthermore, a defining characteristic of CLCA is its reliance on marginal data and system expansion to account for avoided impacts and substitutions (Ekvall & Weidema, 2004). CLCA is able to capture how variations in demand trigger adjustments in production capacity, trade flows, and technology adoption (Guinée et al., 2018). This distinction is particularly relevant in end-of-life (EoL) modelling, where system expansion is widely recognized as the preferred approach for resolving multifunctionality and capturing the environmental benefits of reuse, recycling, and disposal strategies (Schrijvers et al., 2020; Tanguay et al., 2021). By incorporating additional life cycle processes and subtracting avoided burdens, CLCA provides a more realistic assessment of circular economy strategies, which can significantly alter material demand and production patterns.

CLCA offers a valuable means of assessing potential environmental consequences, however, it also presents several methodological challenges. These challenges often stem from system boundary assumptions, data uncertainty, and the complexity of modelling economic interactions. One of the primary difficulties in CLCA lies in its dependence on forecasting economic and industrial dynamics, which introduces significant uncertainty. Van der Voet et al. (2019) emphasize that identifying marginal suppliers in CLCA is sensitive to assumptions about future industrial growth and technological adoption, which may not align with actual developments. Furthermore, CLCA's effectiveness hinges on the availability of high-quality market data, which is often scarce in emerging industries or regions with limited transparency in production and trade flows. These challenges introduce complexities but, CLCA is able to offer a more realistic

assessment of potential environmental consequences due to decision making than ALCA, making it a valuable tool for informing policy and investment decisions.

While ALCA and CLCA share the same core modelling principles, what differentiates them is the selection of processes included in the system model. In CLCA, system boundaries are extended to include processes affected by market dynamics. However, the identification of these processes is often conducted inconsistently, relying on differing assumptions and justifications, which can lead to divergent results. To address this, several researchers recommend the integration of scenario modelling and temporal considerations to support CLCA by providing a more transparent and structured basis for capturing plausible product- and market-related futures, including technological shifts and policy-driven transitions (Zamagni et al., 2012; Heeren & Hellweg, 2019).

Although this thesis employs a consequential framework, this does not imply that one LCA approach is universally superior. Rather, the choice is based on its relevance to the specific questions being addressed. This sentiment is mirrored by Zamagni et al. (2012) that concludes that both ALCA and CLCA have utility and that the distinction between the two lies in how system boundaries are defined—a difference that ultimately depends on a clear and unambiguous articulation of the study's goal.

To further contextualize the distinctions, overlaps, and applications of ALCA and CLCA, readers are encouraged to refer to Ekvall et al. (2020) for a deeper exploration of methodological implications and Schaubroeck (2023) for an industry-oriented perspective on their practical implementation.

2.2 Temporal Considerations in Life Cycle Assessments

While system boundaries and allocation procedures are commonly used to distinguish between LCA typologies, an equally critical factor of LCA modelling is the treatment of time-dependent variables. Many LCA studies, particularly those using steady-state Life Cycle Impact Assessment (LCIA) models, assume static market conditions and technological landscapes, despite real-world shifts in energy systems, material efficiency, and policy frameworks (Beloin-Saint-Pierre et al., 2020). This approach simplifies calculations but can misrepresent environmental consequences, as an instantaneous emission may have significantly different effects compared to a gradual release over time. While this simplification facilitates comparisons across products and industries, it introduces inaccuracies in impact category results due to the aggregation of emissions across space and time (Levasseur et al., 2010).

CLCA is often positioned as more temporally sensitive due to its modelling of system-wide adjustments in response to changes in demand (B. Weidema, 2003). However, its dynamism stems from its focus on marginal suppliers and economic causality—not from an inherent integration of time-based modelling. As Ekvall & Weidema (2004) note, while CLCA evaluates long-term environmental consequences, it does not necessarily model when emissions occur unless time is explicitly incorporated. As Sandén and Karlström (2007) argue, failure to integrate temporal dimensions into CLCA, especially in the context of policy or technological transitions, can result in misleading or incomplete conclusions. As a result, temporal considerations in LCA have gained increasing attention, particularly in studies where the evolution of technologies and market conditions significantly influence results.

Arvidsson et al. (2023) propose a framework that disaggregates LCA into three independent but frequently conflated aspects: maturity, causality, and time. While causality defines whether the LCA aims to reflect actual consequences of decisions (e.g., CLCA) or describe the environmental burdens of a product as it exists (e.g., ALCA), maturity refers to the technological readiness of the system being modelled, and time addresses whether the study is retrospective, contemporaneous, or prospective. This conceptual distinction is especially important when evaluating the integration of temporal dynamics, which are not inherently tied to any one LCA type or level of causality.

This conceptual distinction has clarified the roles of several LCA approaches that explicitly address temporal variation, namely, prospective LCA (pLCA), retrospective LCA (rLCA), and dynamic LCA (d-LCA), as articulated by Arvidsson et al. (2023). Each of these approaches incorporates time in a distinct way and represents a departure from the standard, time-agnostic contemporary LCA. pLCA integrates anticipated changes in technologies, policies, and market structures to model future environmental consequences. rLCA, by contrast, focuses on historical or existing systems and draws from observed data to assess impacts based on past decisions or baseline conditions. d-LCA, meanwhile, captures the timing of emissions and environmental impacts over defined periods, allowing for more accurate modelling of systems where emissions are temporally distributed or delayed (Levasseur et al., 2010; Peñaloza et al., 2016).

In some cases, temporal framing in LCA has been reduced to a binary of retrospective versus prospective analysis, with ALCA commonly associated with the former and CLCA with the latter. This is an oversimplification that overlooks the distinct roles of time and causality in system modelling. Early contributions to this framing include Tillman (2000) and Weidema (2003), who both emphasized that methodological choices should be informed by the decision context. Tillman distinguished between retrospective and prospective LCA approaches based on whether the goal was environmental accounting or decision support, while Weidema characterized ALCA as retrospective, highlighting its utility for hot-spot identification, and CLCA as forward-looking, intended to assess the broader consequences of demand-driven system changes. While this framing has helped clarify the different roles of each approach, it does not imply that ALCA is inherently limited to past-focused assessments or that CLCA must be exclusively prospective.

Building on these early insights, a growing body of literature has cautioned against equating retrospective and prospective framing with attributional and consequential models. Ekvall et al. (2005), Curran et al. (2005), Sandén & Karlström (2007), Zamagni et al. (2012), Buyle et al. (2019), and Arvidsson et al. (2023) argue that time and causality should be treated as analytically distinct dimensions. Both ALCA and CLCA can be applied retrospectively, contemporaneously, or prospectively, depending on system boundaries, modelling assumptions, and the goals of the study.

Despite growing academic consensus on the importance of temporal modelling in LCA, existing international standards offer limited methodological guidance. As noted by Anand & Amor (2017) and Lueddeckens et al. (2020), the ISO 14040 standard acknowledges this limitation, stating that "*the lack of spatial and temporal dimensions in the LCI results introduces uncertainty in the LCIA results. The uncertainty varies with the spatial and temporal characteristics of each impact category*" (ISO, 2006 as cited in Anand & Amor, 2017; Lueddeckens et al., 2020). This raises an essential methodological question: how can LCA

frameworks move beyond static modelling to consistently integrate time in a way that supports forward-looking decision-making?

2.2.1 Prospective LCA

Among the available approaches, pLCA has gained traction for its ability to estimate environmental impacts at discrete future points in time, rather than modelling continuous emissions and impacts as in d-LCA. This distinction is particularly important for systems with extended gaps between life cycle stages, such as those found in infrastructure and durable goods. In these contexts, pLCA enables practitioners to align scenario assumptions with specific intervention points, such as initial production and eventual dismantling or repurposing. As articulated by Arvidsson et al. (2018) and Beloin-Saint-Pierre et al. (2020), pLCA involves a range of modelling strategies, including scenario-based modelling, technology learning curves, and agent- or activity-based approaches, that simulate system evolution through explicit assumptions about the future. These assumptions can apply to both the technosphere and the ecosphere and are typically structured to reflect potential shifts in policy, infrastructure, or industrial practices.

Recent literature highlights the breadth of pLCA applications across sectors where long-term decision-making is required. Levasseur et al. (2010) emphasize the method's relevance to forward-looking planning in energy, mobility, and technology development contexts. In the built environment, for instance, Fouquet et al. (2015) review the use of prospective modelling in assessing buildings and renewable energy systems, where anticipated changes in electricity mix or production processes are integrated at defined future milestones or through continuous time horizons. Similarly, Heeren and Hellweg (2019) apply probabilistic scenario modelling to assess how different urban development trajectories influence material stocks and environmental outcomes through 2100. Their framework varies key drivers such as building turnover rates, per-capita floor area, and material substitution strategies at distinct points in time across multiple scenarios. The results demonstrate that the timing and intensity of these variables significantly shape long-term environmental impacts, especially when considering cumulative emissions and material lock-ins over decades. In this way, the study exemplifies how pLCA can be used not only to forecast impacts but also to identify when future interventions may yield the most substantial environmental benefits.

While studies like Heeren and Hellweg (2019) underscore the strategic value of pLCA for long-term planning, recent literature has also highlighted the methodological complexity involved in its application. A central challenge lies in defining future scenarios that realistically capture how industrial supply chains, energy systems, and material efficiencies are expected to evolve. Arvidsson et al. (2018) proposed two key scenario development strategies: predictive scenarios, which rely on expert elicitation and technological learning curves to forecast likely system evolutions; and bounding scenarios, which bracket best- and worst-case outcomes to test the sensitivity of results. They further emphasized the importance of modelling both foreground and background systems, cautioning that static background inventories may obscure long-term changes in energy, infrastructure, or market conditions.

2.3 Technological Considerations in Life Cycle Assessments

Technological maturity plays a critical role in shaping both the relevance and reliability of LCA results. Maturity refers to the development stage of the technologies or systems being modelled,

ranging from speculative innovations at low technology readiness levels to established, commercial processes (Arvidsson et al., 2023). This dimension directly influences data quality, modelling assumptions, and the nature of uncertainties introduced into the assessment. Importantly, differences in technological maturity also affect the purpose and structure of LCA applications, whether the aim is to support innovation-stage design decisions or long-term system-level planning.

2.3.1 Ex-Ante LCA

Responding to the need to capture both foreground and background evolving system dynamics, ex-ante LCA has emerged as a focused but distinct methodological extension of pLCA, developed to evaluate the potential environmental impacts of emerging technologies at low technology readiness levels, prior to commercialization. Initially used interchangeably with the term pLCA (Arvidsson et al., 2018), more recent work has clarified that ex-ante LCA is distinguished by its focus on low technology readiness level (TRL) or manufacturing readiness level (MRL) innovations, evaluated in hypothetical future deployment scenarios (Arvidsson et al., 2023). Whereas broader applications of pLCA are often used to inform long-term infrastructure planning, policy development, or sector-wide transitions, ex-ante LCA is primarily employed to support research and development and innovation-stage decision-making. As a subset of pLCA, it shares foundational modelling strategies such as scenario-based analysis, technological learning curves, and assumptions about system evolution. Table 2-1 clarifies the distinction between pLCA and ex-ante LCA.

Building on this methodological foundation, the practical aim of ex-ante LCA is to inform design and policy decisions before major capital investments are made, identifying environmental hotspots and optimization opportunities upstream in the development process (Buyle, Audenaert, et al., 2019). In this context, Kravchenko et al. (2019) advocate for the use of leading sustainability indicators to proactively identify trade-offs and indicates their particular relevance in circular economy initiatives where environmental and economic objectives may diverge.

Continued research and advocacy for improvements to temporally oriented LCA approaches aside, there still exists persistent methodological gaps within prospective and ex-ante LCA. These gaps not only hinder comparability across studies but also limit the credibility and decision-usefulness of long-term assessments. A central challenge across both conventional and prospective LCA lies in the treatment of uncertainty and variability, which are introduced at nearly every stage of the modelling process — from system boundary selection to scenario construction and data sourcing. In pLCA, these uncertainties are further magnified by the need to anticipate future technologies, behaviours, and systemic changes. To fully understand the implications of these challenges in pLCA, it is essential to first examine the broader sources of uncertainty and variability in LCA as a methodological approach.

***Table 2-1**: Comparison of Prospective LCA (pLCA) and Ex-Ante LCA*

Dimension	Prospective LCA (pLCA)	Ex-ante LCA
Primary Purpose	Compare environmental outcomes under specified future scenarios	Estimate impacts of emerging / low-TRL technologies before full deployment
Typical Scale	Often meso–macro (products, buildings, portfolios, sectors), but flexible	Often micro–meso (unit processes, products, technology pathways)
Decision Context	Policy, long-lived asset planning, roadmaps, option comparison	R&D direction-setting, pilot-to-scale decisions, innovation management
Time Horizon	Usually mid–long term (commonly 2030–2050+)	Near–mid term (pre-commercialization to early adoption)
System Boundary Emphasis	Broad, includes evolving background systems (e.g., grids, fuels, industry)	Narrow, focuses on foreground detail & scale-up assumptions; background often held consistent or scenario-adjusted
Key Strength	Captures how system evolution changes results and rankings	Identifies design levers and risk hotspots when empirical data are limited
Key Limitation	High sensitivity to scenario choice and future market/energy assumptions	High sensitivity to proxies, yields, and scale-up assumptions (data sparsity)

2.4 Characterizing Uncertainty & Variability in LCA

A foundational framework for understanding uncertainty in LCA was dissected by Lloyd and Ries (2007), who identified three primary sources: model uncertainty, scenario uncertainty, and parameter uncertainty (see Figure 2-2). These categories offer a structured lens through which to evaluate how methodological and modelling decisions propagate variability in LCA results. While such uncertainty is present in all forms of LCA, it becomes especially pronounced in prospective and ex-ante applications, where future-oriented systems must be modelled without empirical validation (Mendoza Beltran et al., 2020; Thonemann et al., 2020).

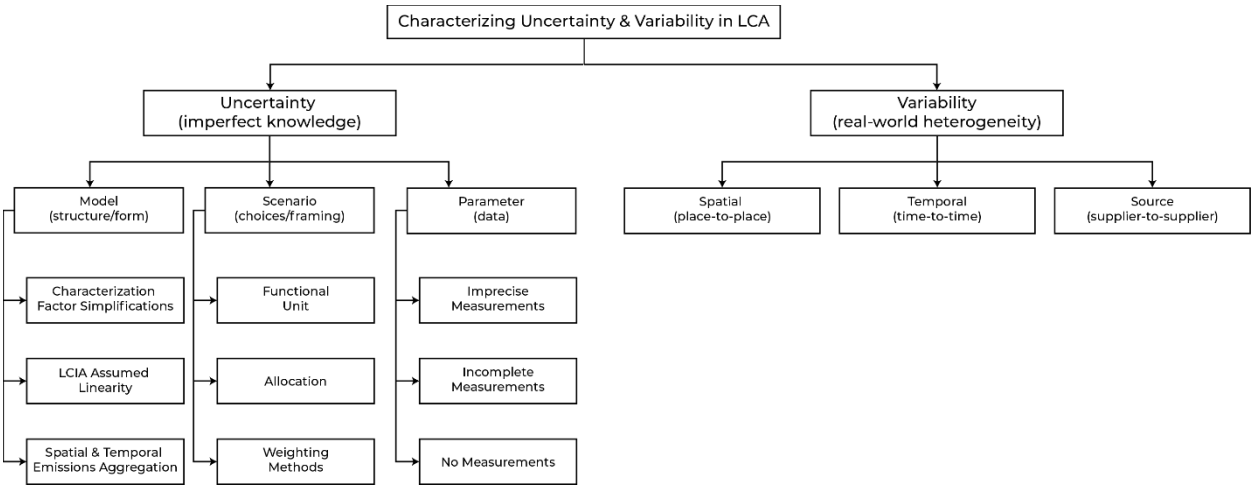


Figure 2-2*: Classification of uncertainty and variability in life cycle assessment (LCA). The diagram separates uncertainty arising from imperfect knowledge or modelling choices—grouped as parameter, scenario, and model uncertainty—from variability, which reflects real-world heterogeneity across space, time, and sources/objects.*

Earlier work by Huijbregts (1998a) complements this typology by introducing a parallel distinction between uncertainty and variability, the latter arising from heterogeneity in environmental systems and datasets. He categorized spatial, temporal, and source-based variability, which stem from intrinsic system differences (e.g., regional emissions profiles or fluctuating resource availability), rather than methodological assumptions. These insights extend the scope of uncertainty analysis beyond purely technical modelling choices to incorporate the natural variability present in real-world systems (Igos et al., 2019).

Beyond categorizing uncertainty by its origin, it is also useful to distinguish uncertainty by its nature. Drawing on Huijbregts (1998a) and formalized by Walker et al. (2003), uncertainty in LCA is commonly divided into epistemic and aleatory forms. Epistemic uncertainty arises from incomplete knowledge and is often reducible. For instance, unclear future policies, unknown technological performance, or poorly justified scenario framing (Björklund, 2002; Igos et al., 2019). Aleatory uncertainty, by contrast, is linked to intrinsic randomness or natural variability in system behaviour, such as weather patterns, demand fluctuations, or material degradation rates. These distinctions are particularly salient in pLCA, where analysts face both: they must assume values for future scenarios (epistemic) while also accounting for inherent system variation (aleatory) (Mendoza Beltran et al., 2020; Van Der Giesen et al., 2020).

This distinction helps clarify why some uncertainties in LCA—particularly in prospective and ex-ante assessments—are so difficult to address: they are not just technical or data-driven but stem from fundamental knowledge gaps about future systems. As such, epistemic uncertainty is particularly relevant in prospective LCA, where assumptions must be made in the absence of empirical validation.

Together, these frameworks highlight that uncertainty in LCA is not monolithic. It varies in source (e.g., modelling choices vs. empirical data gaps), in nature (e.g., epistemic vs. aleatory), and in manifestation (e.g., variability across time or space). This complexity becomes especially

pronounced where analysts must anticipate future system behaviour in the absence of empirical data, thereby amplifying both uncertainty and variability (Thonemann et al., 2020).

In the following sections, each major category of uncertainty and variability is examined in turn, along with how these concerns manifest and are being addressed in the context of contemporary LCA. Particular attention is given to their implications for the credibility, comparability, and robustness of forward-looking assessments.

2.4.1 Model Uncertainty in LCA

Model uncertainty arises from the structural and conceptual assumptions that underlie the modelling framework of a life cycle assessment (Huijbregts et al., 2003). These assumptions include how the system is modelled, which processes are included, and how flows are characterized and aggregated. According to Lloyd and Ries (2007), model uncertainty reflects the simplifications introduced when real-world physical systems are abstracted into manageable analytical constructs. These structural decisions are foundational to any LCA and directly influence its outcomes, yet they are often inconsistently applied across studies or insufficiently reported.

Life cycle impact assessment (LCIA), the phase within LCA where life cycle inventory data—such as emissions and resource uses—are translated into environmental impact indicators (e.g., global warming potential, acidification potential, human toxicity), are particularly sensitive to model uncertainty. Many commonly used LCIA methods apply globally averaged or spatially generic characterization factors (CFs), neglecting regional variability in environmental responses (Bamber et al., 2020). Although regionalized LCIA models have been developed to address this limitation by incorporating spatial differentiation, they frequently rely on generalized or non-region-specific datasets when truly spatially resolved emission inventories are unavailable, introducing additional sources of model uncertainty (C. Mutel et al., 2019). The lack of harmonization across regionalized LCIA approaches limits ecological fidelity and comparability, compounding uncertainty at the impact assessment stage. Furthermore, assumptions such as steady-state behaviour of environmental systems, uniform pollutant mixing, and the exclusion of transformation products contribute additional structural inaccuracies (Huijbregts et al., 2003).

A closely related dimension of LCIA model uncertainty concerns method selection more broadly, including the impact-category set, modelling scope, and temporal framing embedded in a given LCIA framework. In practice, commonly used LCIA methods are not interchangeable: they reflect different modelling choices about which environmental mechanisms are represented, how midpoint and endpoint indicators are constructed, and which conventions are used for spatial and temporal aggregation. These method-level choices are operationalized through the characterization models and associated characterization factors (CFs) provided by the method, which determine how inventory flows are translated into impact indicators and therefore where structural assumptions enter the calculation.

CFs themselves represent a critical source of model uncertainty, as they approximate complex fate, exposure, and effect pathways based on evolving scientific understanding (Finnveden, 2000; Bare, 2010). However, CFs were originally developed as static factors, inherently embedding uncertainty through the simplifications required to represent these intricate environmental processes. Uncertainty reviews also emphasize that LCIA methods should not be treated as neutral “black boxes”; the choice of LCIA framework—including indicator definitions,

reference states, and time horizons where relevant—is repeatedly identified as an important but often under-explored component of model uncertainty relative to the dominant focus on parameter uncertainty (Bamber et al., 2020). As scientific understanding evolves, methodological updates and revisions to underlying characterization models can substantially alter comparative outcomes, particularly for systems where timing, geography, or the handling of multi-compartment pathways is consequential. Accordingly, targeted sensitivity analysis across LCIA methods is increasingly used to test whether conclusions are robust to alternative characterization models and normative choices, especially in contexts involving time-dependent emissions profiles, spatial heterogeneity, or differing interpretations of impact pathways.

Additionally, the development of dynamic CFs where emissions were temporally weighted was introduced to overcome limitations of static CFs (Levasseur et al., 2010; Kendall & Price, 2012). Although improving temporal accuracy, dynamic CFs introduce additional model uncertainty through assumptions about pollutant decay and exposure timing (Lebailly et al., 2014; Igos et al., 2019). In parallel, dynamic life cycle inventories (LCIs) have been developed to capture temporal variations in inventory data, such as changing emission factors and evolving technology efficiencies, thereby improving consistency with dynamic CFs and mitigating temporal model uncertainty (Tiruta-Barna et al., 2016; Sacchi et al., 2022). These shifts illustrate the sensitivity of results to the underlying model structure and parameterization.

As highlighted in a recent review of uncertainty in LCA for the built environment, practitioners often rely on deterministic assumptions and simplified inventory or model structures, which makes transparent documentation of assumptions, adoption of more consistent methodological practices, and use of higher-resolution (spatial/temporal) data critical for improving reliability and interpretability of results (Marsh et al., 2023). Particular attention must be given in pLCA contexts to capture the temporal evolution of systems, clearly communicating the limitations of structural assumptions, and distinguishing between model-based results and scenario-driven influences.

2.4.2 Scenario Uncertainty in LCA

Scenario uncertainty arises from the normative assumptions and value-laden choices made when constructing representations of future conditions (Lloyd & Ries, 2007). Unlike parameter uncertainty, which stems from variability in empirical measurements, scenario uncertainty reflects epistemological limitations. Often these are due to the inability to forecast how the complexities of technologies, markets, policies, or behaviours will evolve over time (Björklund, 2002; Mendoza Beltran et al., 2020). Beltran et al. (2020) state that this form of uncertainty is especially pronounced in prospective and ex-ante LCA, where analysts must simulate systems that have not yet materialized and often rely on constructed futures rather than empirical observations.

Scenario uncertainty is distinct in that it does not arise from imprecision in data, but rather from the subjective decisions embedded within the goal and scope definition of the study. These include assumptions regarding time horizons, technological learning rates, background energy mixes, marginal supplier identification, valuation methods, waste management scenarios, and socioeconomic development pathways (Huijbregts, 1998a; Lloyd & Ries, 2007; Thonemann et al., 2020). Such assumptions serve a dual role: they act as a technical modelling input and simultaneously define the broader narrative within which the product system is evaluated. As

Huijbregts (1998a) describes, these "*uncertainty due to choices*" exert a structural influence over both the inventory and impact assessment of stages of the LCA.

While these assumptions primarily define the future context under which systems are evaluated, their influence extends deeper. Scenario assumptions do not merely affect the interpretation of results; they often reshape the internal structure and temporal behaviour of LCA models themselves. For instance, choosing whether to implement a static or dynamic background system, a decision often treated as a scenario assumption does not change the underlying model structure but fundamentally alters how system dynamics are represented thus, influencing model outputs (Mendoza Beltran et al., 2018). Similarly, assumptions about material substitution pathways or marginal technology displacement, especially prevalent in CLCA, embed forward-looking causal relationships into the internal logic of the model. These examples highlight how scenario assumptions can influence both the external framing and internal mechanics of LCA models, leading to complex interactions between scenario and model uncertainty that are particularly pronounced in forward-looking assessments.

The following subsections outline key modelling assumptions that contribute to scenario uncertainty in prospective and ex-ante LCA. These assumptions—ranging from time horizon selection to socioeconomic projections—not only define future system behaviour but also introduce significant variation into LCA results. While these factors will be revisited in the context of methodological strategies in Section 2.5.2, they are presented here for their role in uncertainty propagation.

2.4.2.1 Sources of Scenario Uncertainty

Scenario uncertainty arises from assumptions about how future systems may evolve, and these assumptions span several core components of an LCA's design. This section disaggregates the primary sources of scenario uncertainty to clarify their individual contributions to uncertainty propagation in LCA. These categories reflect common elements in goal and scope definition that guide model construction and directly influence environmental outcomes.

Time Horizons

The time horizon in LCA defines the temporal scope over which environmental impacts are modelled and interpreted. As outlined by Lueddeckens et al. (2020), time horizons can be grouped into three interrelated categories:

- The assessment time horizon, defined during the goal and scope phase, determines how far into the future system behaviours, interventions, and emissions are projected, establishing the study's temporal scope.
- The inventory time horizon represents the operational lifetime of the product or system and defines the duration over which life cycle processes are tracked, setting the temporal boundaries for data collection and modelling.
- The characterization time horizon governs how emissions are weighted in the impact assessment stage (e.g., 20-, 100-, or 500-year GWP), defining the time frame over which environmental impacts are evaluated.

Scenario uncertainty is most closely tied to the assessment time horizon because it governs how far into the future systems are represented as evolving receiving systems. Because this scope

requires assumptions about future technologies, policies, and behaviours, it introduces epistemic uncertainty. As the assessment horizon extends, these assumptions become less observable and more contingent, widening the plausible outcome space and increasing dependence on scenario architecture. Longer horizons may also require additional structural modelling choices (for example, market boundaries, substitution definitions, and marginal supplier heuristics), which can further widen outcome ranges even under a fixed scenario set.

Short assessment horizons (e.g., 2020–2030) are often selected for their relevance to policy cycles, planning horizons, and actionable decision-making. From an action-oriented perspective, they emphasize the urgency of immediate interventions and make it easier to align LCA studies with short-term regulatory or investment goals. However, they risk underestimating long-term burdens and benefits, particularly for technologies with delayed environmental impacts. These include systems with slow-release emissions (e.g., landfill methane), biogenic carbon storage (e.g., timber sequestration), or extended reuse cycles as part of circular economy strategies (Lueddeckens et al., 2020).

Longer assessment time horizons (e.g., through 2050 or 2100) are often necessary when evaluating technologies expected to contribute to long-term sustainability transitions, such as net-zero carbon pathways or circular economy strategies, because they enable the inclusion of cumulative effects, system feedback, and intergenerational consequences. Such horizons are especially critical in systems involving slow-release emissions, infrastructure turnover, or delayed reuse cycles (Hauschild et al., 2018). However, their extended scope introduces growing epistemic uncertainty, particularly around projections of background decarbonization, technological learning, and policy uptake (Mendoza Beltran et al., 2020).

For instance, Pehnt (2006) demonstrates that long-term reductions in the environmental impacts of photovoltaics and biomass for heating are highly dependent on projected improvements in grid emissions and manufacturing processes. Similarly, Cox et al. (2018) show that electric vehicles outperform conventional vehicles primarily under assumptions of future grid decarbonization. This highlights how conclusions shift based on long-range modelling assumptions. In both cases, the assessment time horizon determines the extent to which these uncertain system dynamics influence outcomes. This illustrates that longer horizons, while necessary for capturing long-term benefits, also introduce compounded epistemic uncertainty and amplify the importance of transparent scenario definition and justification.

In contrast, the inventory and characterization time horizons contribute primarily to structural and model uncertainty, rather than scenario uncertainty. However, in dynamic or scenario-coupled inventories (e.g., evolving electricity mixes during use), inventory modelling can also become a direct pathway through which scenario assumptions propagate into results. Additionally, the inventory horizon can introduce structural uncertainty when relevant post-use processes, such as reuse, recycling, or long-term emissions, fall outside the modelled timeframe. As Lueddeckens et al. (2020) note, such truncation risks excluding flows that may be assumed in interpretation but not represented in the model, potentially distorting system boundaries and comparisons. At the same time, in dynamic or scenario-coupled inventories (for example, electricity supply evolving during the use phase), inventory modelling can also become a direct pathway through which scenario assumptions propagate into results.

The characterization horizon, meanwhile, introduces model uncertainty through its influence on how emissions are temporally weighted during impact assessment. More precisely,

characterization time horizon selection introduces two separable effects: (i) metric-choice sensitivity, reflecting normative temporal weighting (e.g., GWP20 vs GWP100), and (ii) uncertainty embedded in the impact model itself (for example, assumptions governing time-dependent climate response). Levasseur et al. (2010) show that the choice of characterization time horizon and related temporal valuation choices (such as discounting) can materially change whether mitigation options appear to meet climate targets, underscoring how time-horizon selection can shape the perceived viability of mitigation options. Hauschild et al. (2018) further note that applying a fixed characterization window may fail to reflect the temporal lag between emission and impact, particularly in long-lived systems like buildings, where emissions are distributed across extended timeframes. While these latter forms of uncertainty do not arise from scenario assumptions, they can nonetheless introduce inconsistencies that affect model transparency and interpretability.

Figure 2-3 provides a visual separation of the three time-horizon mechanisms and the distinct ways they shape interpretation and uncertainty in LCA.

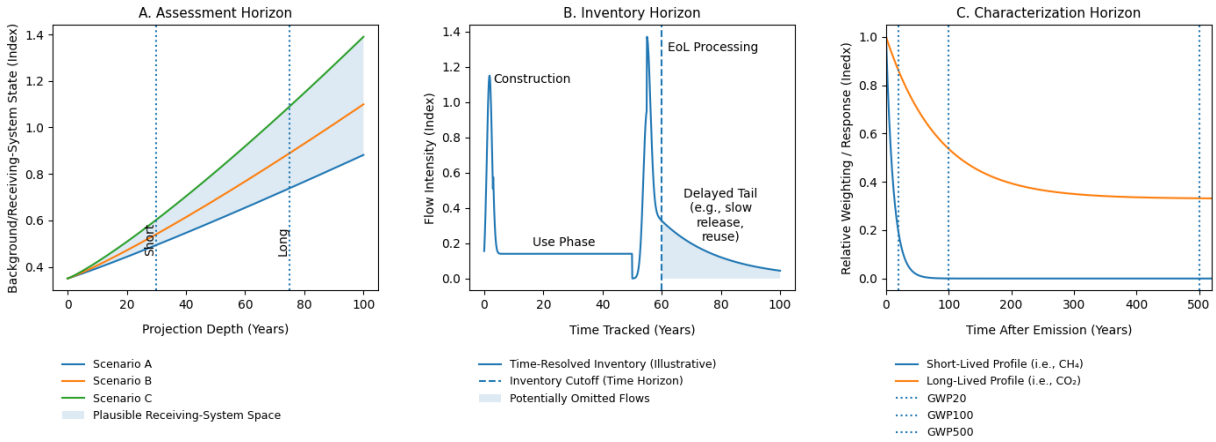


Figure 2-3*: Conceptual schematic distinguishing three LCA time horizons and their dominant uncertainty pathways. (A) The assessment horizon governs projection depth of the receiving system; longer horizons increase dependence on scenario assumptions, illustrated by divergence across plausible receiving-system pathways and an expanding outcome space (schematic, not LCIA results). (B) The inventory horizon defines the time window over which life-cycle flows are tracked; temporally concentrated phases (e.g., construction and end-of-life processing) may appear as pulses relative to the full lifetime, and delayed flows (e.g., slow-release emissions or extended circular loops) can extend beyond the cutoff, creating truncation risk when omitted. (C) The characterization horizon is a metric window length (e.g., GWP20/100/500) applied after emissions occur, influencing temporal weighting and metric sensitivity. Curves and thresholds are illustrative and not to scale.*

Technological Learning and Foreground Dynamics

In pLCA, a central source of scenario uncertainty arises from how technological change is anticipated and modelled within the foreground system. While the evolution of technologies itself is a process of advancement, the act of predicting that evolution becomes a major source of scenario uncertainty. This distinction is critical: technological evolution itself is inevitable, but how we represent it in prospective models, one's expectation or prediction of how it unfolds, is

what introduces scenario uncertainty. This issue is especially acute in ex-ante LCA, where the foreground system includes technologies not yet commercialized, with unknown or non-linear trajectories of improvement, scaling, and adoption.

To address this lack of data, emerging technologies are often assessed based on laboratory-scale experiments, pilot projects, or theoretical efficiencies—data sources that are not only uncertain but also difficult to generalize. Assumptions regarding technology learning curves, scale-up factors, and production efficiencies become key drivers of foreground system behaviour, particularly for novel materials, bio-based products, and energy storage systems. As Arvidsson et al. (2018) note, scenario uncertainty in these contexts stems from both the immaturity of the technology and the methods used to extrapolate future performance.

One widely adopted method for characterizing technological maturity is the use of technology readiness levels (TRLs) and manufacturing readiness levels (MRLs) (Gavankar et al., 2015). TRLs classify technologies on a scale from basic research (TRL 1) to fully operational systems (TRL 9), while MRLs similarly, and subsequently, track the progression of manufacturing capabilities (Buyle, Audenaert, et al., 2019). Lower TRLs and MRLs signal that performance assumptions are more speculative, and that scale-up risks, including unforeseen bottlenecks or non-linear cost and efficiency behaviours, are more pronounced. This tension is captured by the Collingridge dilemma (Collingridge, 1980), which describes how early in a technology's development, the capacity to influence design is high, but reliable knowledge is limited. Later, as systems mature and data improves, opportunities for intervention narrow due to locked-in decisions. pLCA, especially when applied to low-TRL technologies, therefore faces the dual challenge of high uncertainty and high potential influence.

Studies have integrated learning rates derived from historical analogues or sector-specific trends into LCA foreground models, attempting to simulate cost and efficiency improvements over time (Caduff et al., 2012, 2014; Bergerson et al., 2020). However, this approach introduces normative assumptions about innovation trajectories, which are highly sensitive to market dynamics, policy environments, and technical feasibility, often leading to significant uncertainty in projections and potentially biased assessments of future performance and adoption levels.

For example, in studies of lithium-ion battery systems, variation in assumed future energy density or production energy requirements can change the relative climate impact of the technology by more than 40%, depending on which year or version of the process is modelled (Peters et al., 2017). Similarly, bio-based plastic LCA studies have shown that performance metrics such as yield, degradation rates, or co-product treatment assumptions can invert the comparative ranking of materials depending on their scale-up assumptions and baseline conditions (Khoo et al., 2010; Shen et al., 2011).

Background System Dynamics

In LCA, the background system encompasses all processes outside the immediate influence of the decision-maker or technology developer. It includes supporting systems such as electricity generation, material supply chains, transportation networks, and waste management infrastructure are including (Arvidsson et al., 2018). These systems evolve over time due to policy interventions, technological shifts, and structural market dynamics. In prospective and consequential LCA, the way these background changes are represented introduces a distinct

source of scenario uncertainty—uncertainty not about whether change will happen, but about how it will unfold, and how that unfolding is modelled.

While uncertainty in the foreground system often stems from emerging technologies or novel materials, background uncertainty arises from the assumptions made about the evolving system context in which those technologies operate. In pLCA, these assumptions often pertain to future points in time, such as energy mixes or infrastructure in 2030, 2040, or beyond, and reflect efforts to align the background with anticipated implementation timelines (Mendoza Beltran et al., 2020). In contrast, CLCA introduces temporal uncertainty by focusing on system responses over time, identifying marginal suppliers or displaced processes that vary depending on how markets and technologies evolve (Ekvall & Weidema, 2004).

In CLCA, a major source of scenario uncertainty arises from the need to predict how supply systems will react to marginal changes in demand. This includes identifying which suppliers are expected to expand or contract in response to a marginal change, a determination that depends on assumptions about market behaviour, policy incentives, fuel prices, reserve capacity, and technological competitiveness over time. These determinants are both time-sensitive and region-specific, and often involve complex, interdependent variables that are difficult to capture with the granularity available in LCA databases. This introduces a type of scenario uncertainty distinct from characterization or inventory time horizons: rather than reflecting when emissions occur or how long impacts persist, it reflects how supply systems will structurally adjust in response to future decisions. As Ekvall and Weidema (2004) emphasize, different assumptions about future supply elasticity or market competitiveness can lead to different consequential outcomes, even for the same intervention.

For example, Maes et al. (2023) compared different heuristic approaches for identifying marginal electricity suppliers in Western Europe in 2040. Using future energy scenarios derived from integrated assessment models (IAMs)—which simulate long-term interactions between energy, economy, and climate policy—they showed that the selected marginal mix varied significantly depending on the chosen method and time interval. In some cases, the marginal mix was composed almost entirely of renewables; in others, it included a notable share of coal-fired generation. These differences led to substantial variation in global warming potential (GWP), underscoring how methodological assumptions about marginality can dramatically influence the outcomes of consequential LCA.

In pLCA, background system dynamics are also central, but the uncertainty stems from efforts to anticipate how broader systems will evolve over the timeline of a technology's implementation. For instance, Arvidsson et al. (2015) showed that the climate impact of cellulose nanofibril (CNF) production via mechanical treatment was highly sensitive to the assumed electricity mix. When shifting from Sweden’s low-carbon mix to a coal-dominated Polish grid, GWP increased by up to 1800% for some CNF process variants. This underscores how evolving background conditions, especially energy systems, can dominate life cycle impacts in energy-intensive technologies.

As pLCAs increasingly integrate consequential modelling to capture indirect effects, assumptions about background dynamics are often layered across both frameworks. This “stacking” of assumptions heightens the need for transparency, as even small shifts in projected infrastructure or supply systems can invert comparative conclusions between technologies. A further source of structural uncertainty in these hybrid applications concerns response timing in

marginal supplier identification. Even when the analyst holds a single future background trajectory fixed, marginal mixes can differ depending on whether marginal supply is interpreted as a time-slice snapshot at the assessment year (i.e., capacity available in that year) or as a delayed response where investment and commissioning occur after technology- and sector-specific lead times. In practice, this choice can shift the effective start and end years of the trend window used to detect expanding suppliers and therefore alter which technologies are identified as marginal. More broadly, temporal misalignment between background assumptions and foreground deployment schedules can further exacerbate scenario uncertainty. For example, modelling a 2040 technology using a 2025 background may understate its environmental performance, while applying a 2080 background may unrealistically inflate avoided burdens. Together, these issues reinforce that prospective consequential results depend not only on projected background trajectories, but also on explicit temporal assumptions about when technologies operate and how quickly supply systems respond to demand signals (e.g., Maes et al., 2023).

To address this challenge, analysts have proposed strategies such as temporal harmonization—aligning foreground deployment timelines with background assumptions—and scenario layering, where multiple time- and region-specific scenarios are tested in parallel to assess robustness. IAMs and Shared Socioeconomic Pathways (SSPs) are frequently used to structure these background transitions and support alignment. As Arvidsson et al. (2018) note, failure to reflect future contextual changes reduces the credibility and relevance of pLCA outcomes, making temporal consistency between systems a core requirement for forward-looking assessments.

2.4.2.2 Socioeconomic and Policy Pathways

Understanding how socioeconomic and policy developments influence environmental outcomes is central to scenario design in pLCA. These trajectories are particularly salient in pLCA due to their explicit temporal framing and influence over long-term system evolution. Although causal relationships and system boundary choices in ALCA and CLCA can influence how such uncertainties manifest, the integration of these pathways in pLCA is fundamentally driven by assumptions about how society, institutions, and policies will co-evolve over extended time horizons. These assumptions introduce scenario uncertainty, particularly when modelling complex systems over multiple decades (Gütschow et al., 2021).

Shared Socioeconomic Pathways (SSPs)

A common entry point for structuring such futures is the Shared Socioeconomic Pathways (SSPs) which are a set of five standardized, narrative-driven scenarios developed to represent plausible global futures over the 21st century (O'Neill et al., 2014). Each SSP describes a distinct combination of trends in population growth, education, urbanization, technological change, institutional strength, and climate policy (Kriegler et al., 2012). For example, SSP1 represents a world characterized by sustainability and coordinated climate action, while SSP5 depicts a future of fossil-fuelled development and rapid economic growth with limited environmental regulation.

These narratives were designed not as predictions, but as boundary conditions to explore how social and institutional factors influence mitigation and adaptation capacity. In themselves, SSPs are qualitative, they must be operationalized through models to assess their implications for emissions, infrastructure, and resource use. This task is carried out using Integrated Assessment Models.

Integrated Assessment Models (IAMs)

Integrated Assessment Models (IAMs) are dynamic modelling frameworks that simulate the interactions between human and environmental systems. They translate SSP storylines into quantitative projections of energy use, land dynamics, industrial systems, emissions, and economic activity (Moss et al., 2010; Harmsen et al., 2021). IAMs are widely used in climate mitigation research to explore how global and regional systems might evolve under various policy, technology, and behavioural assumptions. In the context of pLCA, they provide scenario-consistent projections used to align both foreground technologies and background inventories with plausible future system states.

IAMs vary in scope and structure, but in pLCA, most studies rely on process-based IAMs—models that simulate energy systems, land use, and macroeconomic development in sectoral and regional detail (Mendoza Beltran et al., 2020; Maes et al., 2021; Sacchi et al., 2022; Maes et al., 2023, 2024). These are distinct from more abstract cost–benefit IAMs (e.g., DICE, FUND), which focus on optimizing global climate policy based on economic damage functions and are not commonly used in life cycle modelling. Notably, process-based IAMs underpin the IPCC's transformation pathway analyses, reflecting their detailed representation of technological, macroeconomic, and biophysical systems (Van Beek et al., 2020; Harmsen et al., 2021).

Among process-based models, IMAGE and REMIND are the most frequently applied in pLCA studies due to their integration with SSPs, sectoral coverage, and the availability of publicly accessible, scenario-aligned outputs that can be mapped to life cycle inventories. Their outputs have been used extensively in the development of prospective life cycle inventory (LCI) databases, particularly through tools such as premise, which translates scenario outputs into modified ecoinvent datasets (Sacchi et al., 2022; Mendoza Beltran et al., 2020; Šimaitis et al., 2024). Other IAMs, including GCAM, MESSAGE-GLOBIOM, and AIM, also contribute to SSP quantification and have been applied in LCA frameworks, though they appear less frequently in pLCA literature. The integration of SSPs and IAMs is depicted in Figure 2-4.

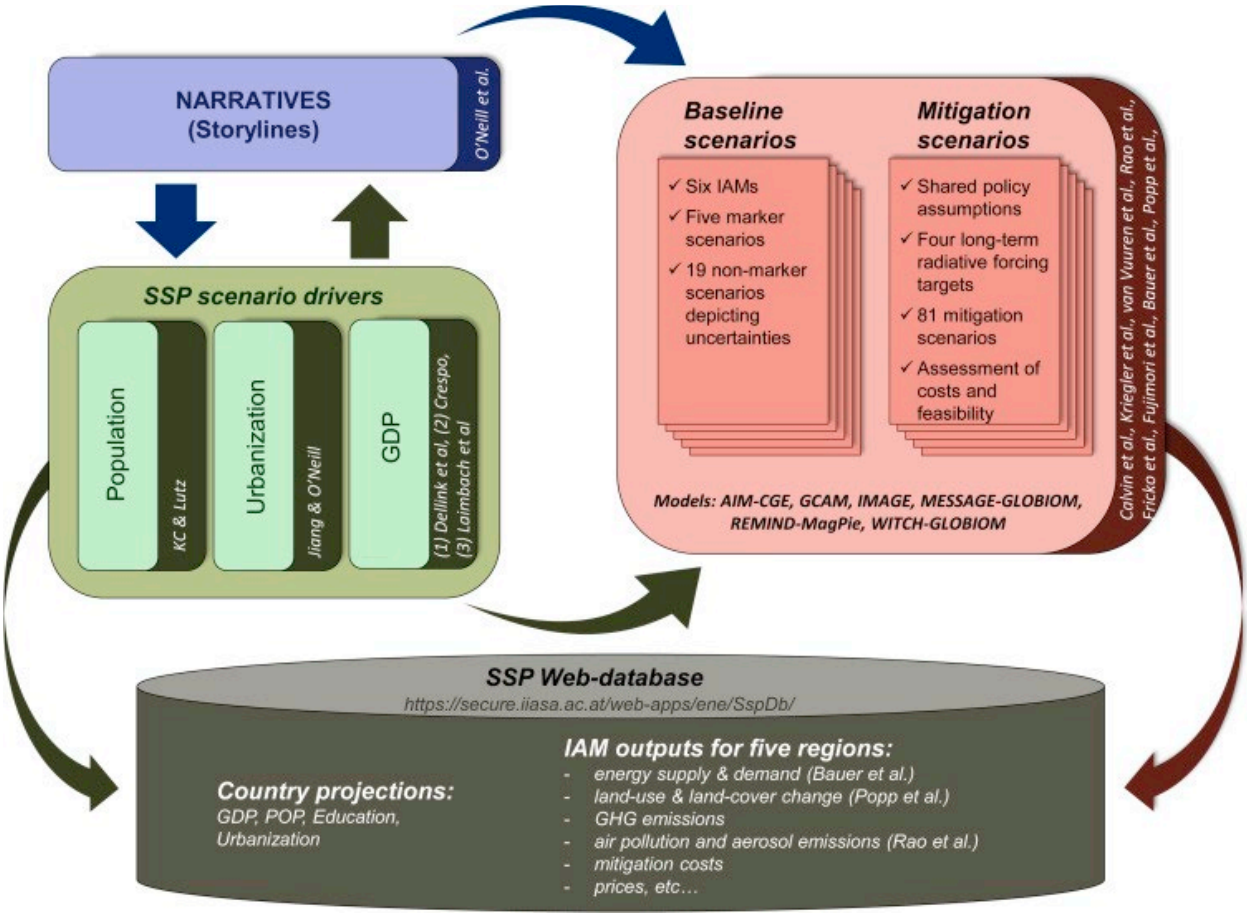


Figure 2-4*: The sequence of information used to project future levels of climate change (derived from Riahi et al., 2017)*

Scenario Outcomes in IAMs

In many IAM-based scenario outputs, pathways require a concise label for the outcome stringency, or in other words, how hard the world mitigates. Different IAMs (IMAGE, REMIND, etc.) expose that stringency with different conventions:

- ***IMAGE***: Labels reflect radiative- forcing stringency characterized as H/M/L/VL for high/medium/low/very-low forcing, sometimes with LO/HO to denote low/high overshoot (e.g., SSP2-VLLO = very-low forcing, low overshoot) (Sacchi et al., 2018/2023).
- ***REMIND***: Uses policy constructs such as Nationally Determined Contributions (NDC) (UNFCCC, 2025), National Policies Implemented (NPi), and explicit temperature/peak-budget constraints.
- ***TIAM-UCL***: Often tagged with Representative Concentration Pathways (RCPs), which correspond to different levels of radiative forcing expected by the year 2100 (e.g., 2.6, 4.5, or 8.5 W/m²).

While these labels are often used in combination with Shared Socioeconomic Pathways (e.g., SSP1–2.6 or SSP5-H), they denote climate outcomes, radiative forcing or temperature targets, rather than socioeconomic drivers. They originated to standardize climate model experiments

and do not, by themselves, encode assumptions about demographics, technology, or policy. In pLCA (Moss et al., 2010; Van Beek et al., 2020), these outcome labels are best treated as concise indicators of mitigation stringency; the technological and economic pathways that shape prospective LCIs come from the IAM runs that meet those outcomes. Figure 2-4 depicts the sequence of information to project future levels of climate change. Unless a study is explicitly evaluating climate targets, the labels serve a descriptive rather than operational role. Importantly however, different IAMs can achieve the same outcome target via different sectoral pathways, which introduces meaningful scenario uncertainty for pLCA.

IAM-Specific Uncertainty

Even when IAMs are aligned to the same SSP narrative, significant scenario uncertainty arises from how different models interpret and operationalize those shared assumptions. Integrated assessment models vary in their representation of technological change, energy system architecture, land-use strategies, and policy mechanisms (Intergovernmental Panel On Climate Change (IPCC), 2023). These structural differences lead to divergent projections for emissions intensities, electricity mixes, material flows, and land footprints, each of which directly affects background inventories used in pLCA.

For example, Georgiades et al. (2023) used outputs from the IMAGE model under an SSP2-based scenario to project cement-sector emissions in Europe. Their results showed that substituting Portland cement clinker with supplementary cementitious materials (SCMs) could reduce climate change impacts by approximately 42% by 2050 relative to 2020, with reduction potentials reaching up to 88% under high clinker replacement scenarios reported in the literature. These findings illustrate how the scale of mitigation achievable through material substitution can vary substantially depending on process assumptions and background conditions derived from IAM scenarios, highlighting the sensitivity of prospective LCA results to underlying scenario frameworks.

Similarly, Maes et al. (2024) demonstrated that pLCA outcomes for cement-sector mitigation are highly sensitive to the structural features of the IAM used to supply background data. In their case, clinker production was linked to the IMAGE IAM, which uniquely models cement demand down to the clinker level and includes investment choices for kiln technologies. However, the authors identified multiple model limitations, such as the overrepresentation of natural gas in the fuel mix and the absence of alternative kiln technologies or novel binders, that constrained the range of mitigation options represented.

These examples underscore that IAM selection is a key driver of scenario uncertainty in pLCA. Although multiple IAMs may implement the same SSP framework, their internal modelling logics can yield markedly different environmental outlooks. This uncertainty is particularly consequential when IAM outputs are used to support prospective comparisons or policy decisions in LCA.

2.4.2.3 Interdependencies and Structural Scenario Design

Scenario uncertainty in pLCA does not arise in isolation. Rather, it is often entangled with model structure, temporal assumptions, and causality frameworks especially in studies using consequential logic. This section unpacks how these interdependencies shape and amplify scenario uncertainty, particularly through modelling logic that embeds future expectations and normative judgments.

Causality and Modelling Logic

The way scenario uncertainty is addressed in life cycle assessment is closely tied to the structural logic of the model, specifically, how causality, time, and system boundaries are framed. This interdependence becomes especially clear when comparing attributional and consequential approaches. In ALCA, scenario assumptions are typically implicit and limited to methodological choices such as allocation or impact assessment. CLCA, in contrast, embeds scenario logic more deeply by modelling system-wide responses such as marginal technology shifts or policy-driven market effects making it inherently more sensitive to assumptions about the future (Ekvall & Weidema, 2004).

This distinction is reflected in a review by Bamber et al. (2020) , who found that while 95% of CLCA studies addressed parameter uncertainty, only 32% addressed scenario uncertainty, and just 16% considered model uncertainty. Their findings emphasize that in consequential modelling, assumptions about marginality and substitution are not just inputs—they are embedded in the model's causal structure. Although ALCA is less structurally sensitive to such assumptions, prospective applications still introduce scenario uncertainty, particularly when used to assess future technologies or infrastructure. Recognizing how these structural interdependencies affect uncertainty propagation is essential for interpreting results across LCA approaches.

Interdependencies Between Assumptions

In LCA, and more specifically pLCA, modelling choices often operate in tandem, meaning that the influence of one assumption depends on others. For example, a decision to adopt a long-term assessment horizon may implicitly justify a dynamic background system, which in turn shifts marginal supplier identification and affects allocation strategies. These cascading effects reflect the structural interdependencies described by Thonemann et al. (2020), where overlapping uncertainties across data quality, comparability, and scenario framing collectively shape the trajectory of modelled outcomes.

This has practical implications: altering a scenario input in isolation may not reflect the full influence of a plausible future. For example, assuming a decarbonized grid by 2050 implies changes not only in electricity supply but also in demand, infrastructure, and marginal suppliers. Modelling only the grid change while holding other systems static may lead to oversimplified outcomes. Mendoza Beltran et al. (2020) demonstrate this by limiting changes to the electricity sector in their prospective LCA, while explicitly acknowledging that expanding scenario alignment across sectors would improve robustness and scenario consistency.

Normative Assumptions and the Ethics of Scenario Framing

While earlier sections focused on technical and structural uncertainties, this section turns to the normative assumptions embedded in scenario design, assumptions that shape the model's causal logic. In pLCA, choices like system boundaries, allocation methods, or functional units often reflect expectations about how systems should evolve, not just how they might. For instance, modelling future reuse assumes supportive social, economic, and policy conditions; in doing so, it embeds a normative stance within what might otherwise appear to be a purely technical decision.

This blurring between scenario design and model structure complicates uncertainty treatment and raises concerns about transparency and reproducibility. Cherubini et al. (2018) show how substitution assumptions in system expansion strongly influence outcomes but are rarely tested. Similarly, Mutel et al. (2019) illustrate how spatial aggregation embeds assumptions about regional uniformity, masking variation tied to future narratives. Such conceptual uncertainty cannot be captured through parameter ranges alone and demands more explicit reflection in prospective LCA.

2.4.3 Parameter Uncertainty in LCA

Parameter uncertainty, also referred to as stochastic or data uncertainty (Bamber et al., 2020), arises from imprecision or variability in the numerical values used to populate an LCA model (Huijbregts, 1998a; Lloyd & Ries, 2007). It reflects limitations in the empirical data used to quantify system inputs and outputs, including material flows, energy use, emissions, and characterization factors. Unlike scenario uncertainty, which stems from assumptions about future developments, and model uncertainty, which arises from structural modelling choices, parameter uncertainty is rooted in empirical limitations such as measurement error, missing data, outdated inventories, and lack of spatial or temporal specificity.

This form of uncertainty occurs throughout the LCA framework, particularly in the LCI and impact assessment phases (Reap et al., 2008). In the foreground system, it often stems from site-specific variation, limited access to primary data, or inconsistencies in process measurements. These values are typically estimated or collected by the analyst and can vary significantly across locations or facilities. The challenge is especially acute in ex-ante LCA, where analysts often extrapolate laboratory- or pilot-scale data to industrial conditions, introducing epistemic uncertainty around yields, material intensities, or energy requirements (Thonemann et al., 2020). Because foreground systems are typically closest to the decision-making intervention, uncertainty at this level can disproportionately influence comparative results.

In contrast, background system parameter uncertainty stems from the structure and abstraction of LCI databases. These databases frequently rely on averaged or proxy datasets that may not capture the specific technological, geographic, or temporal context of the modelled system (Marsh et al., 2023). When such data is used without adjustment, it introduces uncertainty through representativeness gaps. Moreover, as Marsh et al. note, the ability to characterize and propagate uncertainty differs across LCA tools and database architectures, resulting in inconsistencies in how parameter uncertainty is modelled and interpreted.

In prospective and ex-ante LCA, these challenges are further amplified. Analysts often rely on theoretical process models, patent data, or analogies to mature technologies, compounding uncertainty in both material and energy flows (Thonemann et al., 2020). Several studies suggest that parameter uncertainty can dominate outcome variability. For example, Huijbregts et al. (2003) found parameter uncertainty to be the most influential factor in a comparative LCA of insulation materials. However, this dominance is context dependent. Steen (1997) and others emphasize that in some cases, model and scenario uncertainties can be more impactful, underscoring the need for case-by-case assessments.

Despite these limitations, parameter uncertainty remains the most addressed form of uncertainty in LCA. Bamber et al. (2020) found that it was considered in 87% of attributional and 95% of consequential LCA studies that included uncertainty analysis, far exceeding the treatment of

scenario or model uncertainty. While the computational effort required remains a barrier (Marsh et al., 2023), most LCA software packages support Monte Carlo simulation, making parameter uncertainty technically more accessible than other uncertainty types and contributing to its widespread uptake.

2.4.3.1 Sources of Parameter Uncertainty

This subsection outlines the primary sources of parameter uncertainty that influence LCA outcomes, particularly in prospective and comparative studies. The sources are grouped according to their origin: issues of data quality and representativeness, modelling assumptions in the foreground system, abstraction and implementation in background datasets, and parameterization of the impact assessment. Each category highlights how uncertainty can propagate through the model and shape result interpretation, with emphasis on examples that illustrate common pitfalls in current LCA practice.

Data Quality

Data quality, as described by Weidema and Wesnæs (1996), refers to the characteristics of data as expressed through its accompanying metadata. These characteristics include uncertainty (distribution and variance), reliability (source and verification), completeness (coverage and sampling), and alignment with the temporal, geographical, and technological scope of the study. As such, representativeness, specifically the alignment of data with relevant contextual dimensions, will be treated separately in the following subsection. This subsection focuses instead on issues related to the intrinsic integrity, traceability, and reliability of the data itself. These characteristics of data quality remain among the most persistent and structurally embedded sources of parameter uncertainty in LCA.

Uncertainty of Data Quality

Uncertainty, in the context of data quality, refers to the extent of numerical imprecision associated with LCI inputs, typically captured through variance, standard deviation, or probability distributions. In theory, this uncertainty should be characterized by assigning statistical ranges to reflect the variability of input data. As Steen (1997) notes, “all data are uncertain to some extent,” and may be described by distributions such as normal or lognormal curves. However, most LCI datasets present only point estimates, without associated variance or confidence intervals. For instance, Steen elucidates that the point value is often an average of historic data which is intended to be applied to a future scenario, often under different circumstances, adding layered uncertainty. This limits the ability of practitioners to propagate uncertainty using probabilistic techniques such as Monte Carlo simulation. Maurice et al. (2000) emphasize that the lack of empirical support for assigning distributions means analysts often rely on expert judgment or wide, assumed ranges.

Numerous efforts have been made to categorize the sources of uncertainty arising from poor data quality. It is difficult to detect and even harder to quantify, particularly when datasets are reported without associated metadata or confidence intervals (Reap et al., 2008). As early as the mid-1990s, Fava et al. (1994) and Weidema and Wesnæs (1996) emphasized that the absence of harmonized data collection practices undermines result comparability. They advocated for increased standardization through the use of data quality assessments (DQAs) and the pedigree matrix, respectively. Likewise, Vigon & Jensen (1995) stressed that the way data is collected, compiled, and presented has direct implications for reliability and interpretability. In many cases,

inventory data originates from sources with unclear measurement protocols, outdate reference periods, or inconsistent system boundaries, compounding parameter uncertainty across both foreground and background systems.

The lack of formal uncertainty characterization becomes particularly problematic in comparative LCA, where modelled differences between alternatives may fall within overlapping uncertainty bounds. Yet without explicit ranges, these overlaps remain untested, and differences may be overstated. This challenge is compounded by the default assumption of parameter independence in most LCA software. As Björklund (2002) notes, assuming input parameters are statistically independent, when in fact they may be correlated, can inflate modelled output uncertainty and does not reflect real-world interdependencies.

Maurice et al. (2000) illustrate this concern in their study of fossil-fuel power plants, where emissions such as SO_2 were shown to be strongly correlated with energy output. They found that ignoring these dependencies, and treating parameters as statistically independent, distorted the propagation of uncertainty across the model. This example highlights the importance of recognizing interdependencies among process parameters, and the risk of oversimplification when such relationships are excluded from uncertainty analysis. As Steen (1997) observed, "the estimation of uncertainty in itself is very uncertain in LCA studies," which may explain its frequent omission in practice despite its foundational role in interpretability, comparability, and decision-making relevance.

These concerns are further amplified in both prospective and consequential LCA applications. In pLCA, historical data used to model emerging technologies may be poorly suited for forecasting future system behaviour, particularly when extrapolated from low-TRL processes. Meanwhile, CLCA depends heavily on short-term market responses and marginal technology modelling, which requires timely, high-resolution data to identify displaced suppliers or estimate avoided burdens. While these are primarily expressions of scenario uncertainty—arising from assumptions about how future systems may behave—the underlying data used to operationalize these assumptions often carry significant parameter uncertainty. This dual-layered uncertainty, where imprecise or poorly characterized inputs are embedded within speculative futures, complicates interpretation and reinforces the need for transparency in both modelling logic and data quality. Beyond questions of numerical uncertainty, another key dimension of data quality is the reliability of the data source itself.

Reliability of Data Quality

Reliability in life cycle assessment refers to the level of confidence in the origin, traceability, and methodological transparency of input data. It concerns the credibility of data sources, including whether the data are directly measured, modelled based on analogues, or assumed without empirical validation. Essentially, reliability relates not only to the completeness of data flows but also to the trustworthiness of the data present, whether it is measured or modelled, disclosed or assumed, and verified or undocumented. Weidema and Wesnæs (1996) and Vigon & Jensen (1995) both argue that unclear sourcing, lack of verification, and undocumented assumptions reduce confidence in LCI results.

In practice, data reliability issues frequently stem from reliance on unpublished internal studies, generic proxies, or secondary databases with unknown assumptions or outdated methods. For instance, Marsh et al. (2023) report that many environmental product declarations (EPDs) omit

required information, apply inconsistent allocation procedures, and fail to specify key assumptions, leading to results that are neither fully transparent nor easily comparable. In one cited review of 436 EPDs, only 0.04% were considered comparable, and just 5.04% contained all mandatory information, directly implicating poor documentation and reliability in widespread modelling challenges.

Critically, while reliability relates to the traceability and credibility of the data included, it often goes hand in hand with completeness, which concerns whether all relevant flows have been included at all. As the following section will illustrate, omissions in system coverage can be equally consequential for uncertainty propagation in LCA models.

Completeness of Data Quality

Unlike reliability, which pertains to the credibility of available data, completeness addresses whether those data exist at all. Omissions can arise from cutoff thresholds, scope limitations, undocumented assumptions, or failure to track auxiliary flows, all of which introduce parameter uncertainty by silently removing potential contributors to overall system performance.

Ayres (1995) illustrates this issue in his critique of two European LCAs (i.e., FOEFL and PWMI), where emissions were reported in aggregated categories such as “hydrocarbons,” omitting toxic substances like benzene or styrene. Additionally, pollution control measures were not specified, leaving key treatment variables unreported. These exclusions obscured meaningful differences between alternatives and reflected a level of system truncation that compromised the interpretability of results.

In background systems, completeness challenges often stem from the structure of secondary databases and standardized inventories, which are designed for consistency and usability rather than exhaustive system coverage. As discussed in the previous section on reliability, many EPDs and database entries omit or aggregate flows in ways that obscure system-level distinctions. Björklund (2002) warns that such generalization can create a false sense of precision when variation exists but is not visible to the modeler.

In foreground systems, completeness gaps more often stem from missing data on auxiliary materials, co-products, or untracked by-products. As Reap et al. (2008) observe, databases rarely flag records as incomplete, leading analysts to assume full system coverage. These issues are especially problematic in comparative LCAs, where differences in inventory scope can distort conclusions even when core system functions appear similar.

While closely related, completeness is not synonymous with reliability or representativeness. A dataset may include all relevant flows but still rely on unverifiable or outdated sources, compromising reliability. Likewise, a complete and reliable dataset may fail to reflect the correct geography, time frame, or technological context—raising representativeness concerns. These interdependencies highlight how parameter uncertainty can propagate through different facets of data quality, a relationship examined further in the next section.

Representativeness

Representativeness refers to the extent to which input data reflects the specific technological, temporal, and geographic context of the system being modelled. Even when data appears precise or complete, it may introduce parameter uncertainty if it does not accurately correspond to the conditions of the study. As noted in multiple LCA reviews (Reap et al., 2008; Marsh et al.,

2023), issues of poor representativeness are among the most common and consequential contributors to uncertainty, particularly in comparative or prospective applications.

Temporal mismatch occurs when data collected in the past is used to represent current or future systems. This is especially problematic in background systems such as energy supply, where decarbonization trends can rapidly alter the emissions intensity of electricity generation. For example, using 2010 grid mix data to model present-day or 2030 scenarios may underestimate the benefits of electrification in technologies like electric vehicles or heat pumps. Such temporal misalignment can skew absolute results and distort relative rankings between alternatives. LCA studies that model future scenarios, such as Pehnt (2006) or Mendoza Beltran et al. (2020), demonstrate that even modest shifts in energy infrastructure can produce large differences in global warming potential.

Geographic mismatch arises when data from one region is used to model a system in another without appropriate adjustment. This is a common issue in global LCAs that use datasets from European contexts (e.g., ecoinvent or GaBi) to represent production systems in Asia, North America, or Africa (Wernet et al., 2016). While ecoinvent v3 introduced geographically explicit markets to address this limitation, coverage remains uneven across sectors, and Rest-of-World (RoW) datasets are still frequently used as proxies. As Steubing et al. (2016) note, this can lead to significant misrepresentation of impacts, particularly when region-specific activities are missing or generalized from dissimilar contexts.

Technological mismatch is another frequent source of representativeness uncertainty, especially in assessments of emerging technologies. When primary data are unavailable, analysts often rely on proxy data from similar but not identical technologies. As Milà i Canals et al. (2011) categorize, proxy data can take several forms, including direct proxies (using a single surrogate product), averaged proxies (using an average of multiple similar products), and scaled proxies (linearly scaling existing data to approximate missing parts). The authors note that while direct proxies are the simplest to apply, they introduce significant uncertainty due to intrinsic differences between the proxy and target systems, such as variations in cultivation practices, yields, or process inputs. Scaled proxies offer a method to reduce this uncertainty by adapting data using known system parameters, though this requires detailed knowledge and can still leave residual variance.

Such mismatches result in data that are not truly representative of the technology under study, thereby introducing uncertainty that may distort comparative conclusions (Thonemann et al., 2020). These issues are compounded by limited documentation on the original data context in many background databases, which often present averaged or typical inventories without sufficient metadata. Consequently, practitioners may unknowingly apply data that are poorly aligned with their modelled system, creating silent but significant uncertainty that complicates result interpretation.

Scaling & Extrapolation

In prospective and ex-ante LCAs, a key source of parameter uncertainty arises from the derivation and adaptation of input data through scaling and extrapolation, particularly when projecting emerging technologies from laboratory or pilot scales to commercial-scale systems. These data transformations introduce variability in key parameters such as yields, energy consumption, and emissions, which are often imprecisely known due to limited empirical

observations at early technology readiness levels (TRLs) (Arvidsson et al., 2018; Moni et al., 2020).

While the selection of scaling pathways, such as assumptions about technology learning curves, adoption rates, and system integration, is a matter of scenario uncertainty (discussed in Section 2.3.2), the parameter uncertainty lies in the inherent variability and quality of the data used to model these projections. For instance, approximating future energy demand and fuel mixes within decarbonization scenarios involves parameter uncertainty in emission factors and operational assumptions used to represent these systems (Mendoza Beltran et al., 2020; Georgiades et al., 2023).

Furthermore, user behaviour and operational variability contribute to parameter uncertainty, as assumptions on system run-time, maintenance, or degradation over a product's lifespan are typically based on limited or idealized data rather than comprehensive real-world measurements. For example, Peters et al. (2017) show that assumptions about lithium-ion battery lifespan and use intensity significantly influence their environmental performance in LCA studies.

Because scaling and extrapolation are often performed manually by practitioners to address data gaps, the assumptions involved may not be fully documented or critically evaluated. This lack of transparency can lead to underestimation of parameter uncertainty. Figure 2-5 links this issue to the prospective-LCA timing problem: at the assessment time t_0, emerging technologies are often at low–mid TRL and empirical evidence is limited while design freedom is high, so scaled parameters (e.g., yields and energy intensity) span wide plausible ranges. As adoption and operational evidence accumulate toward t_f, parameter ranges typically narrow, even as opportunities to materially redesign the system diminish.

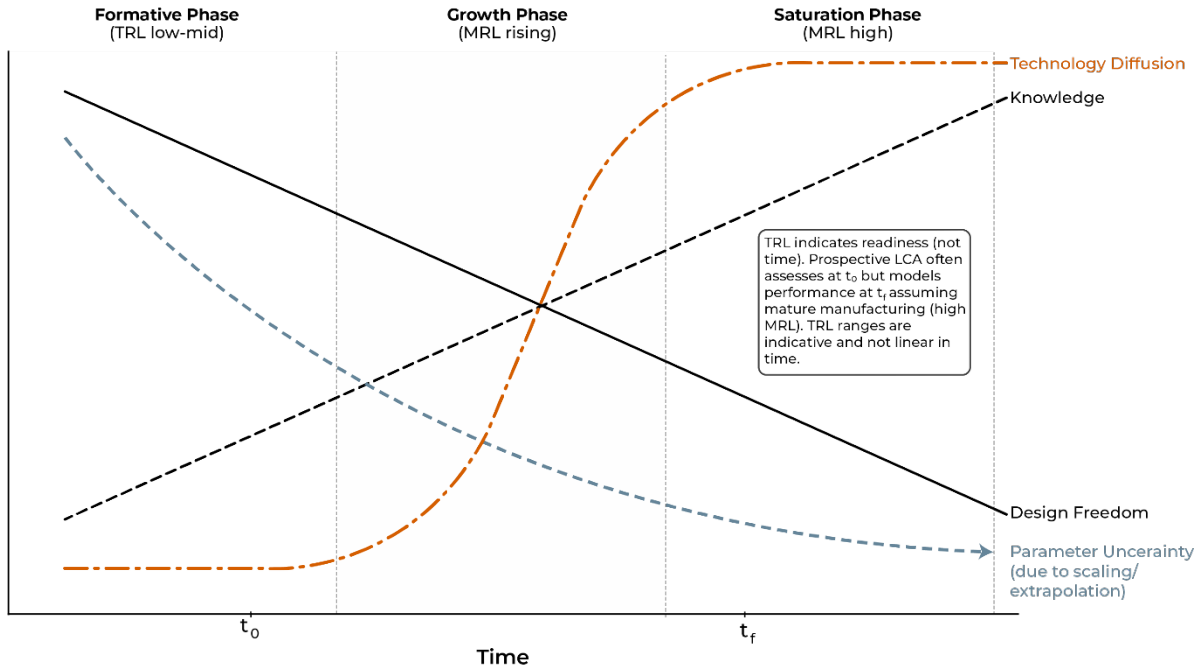


Figure 2-5: *Illustrative graphic of the relationship between technology diffusion, knowledge accumulation, and design degrees of freedom across time in prospective LCA, with an indicative trajectory of parameter uncertainty due to scaling/extrapolation. TRL/MRL is shown as a*

readiness concept relevant at the assessment time t_0, while modelled performance at t_f typically assumes manufacturing readiness (adapted from Arvidsson et al., 2018).

2.4.4 Variability in LCA

Environmental LCAs aim to support decision-making that leads to improved environmental outcomes. However, ignoring the inherent variability in environmental systems and processes risks misleading conclusions. Unlike uncertainty—which arises from lack of knowledge or data—variability reflects real, inherent heterogeneity in environmental conditions, technological performance, or operational contexts that cannot be reduced by better measurement (Huijbregts, 1998a; Marsh et al., 2023).

Heijungs and Huijbregts (2004) as well as Lloyd and Ries (2007) classify variability into three primary types:

- ***Spatial Variability:*** Differences in environmental conditions, emissions, and impacts across geographic locations. For example, acidification impacts depend heavily on local soil buffering capacity, meaning global or regional average characterization factors may misrepresent true environmental effects (Huijbregts, 1998a; Milà i Canals et al., 2011).
- ***Temporal Variability:*** Changes in system behaviour, emissions, or environmental responses over time. Hourly fluctuations in electricity supply mixes can shift environmental impacts by over 200% compared to annual averages, especially in energy-intensive sectors like power and transport (Amor et al., 2014).
- ***Source-Based Variability:*** Differences arising from heterogeneity among data sources, technological variants, or operational practices. For instance, steel production can vary significantly in energy use and emissions depending on plant efficiency, region, and scale—a variability often masked by aggregated LCI datasets (Marsh et al., 2023).

These types manifest in both life cycle inventory (LCI) data, such as emission factors or energy mixes, and life cycle impact assessment (LCIA) models, including fate and exposure processes (Huijbregts, 1998a; Hauschild et al., 2018). Although variability cannot be reduced through improved data or practices, accurately characterizing and acknowledging it improves LCA by enabling more informed interpretation of results.

2.4.5 Interplay Between Variability and Uncertainty in LCA

Variability and uncertainty in life cycle assessment are inherently intertwined and mutually influential, shaping the interpretation and reliability of LCA results. Variability, reflecting real-world heterogeneity in environmental conditions, processes, and technologies, exists alongside multiple forms of uncertainty, including parameter uncertainty (due to limited or imprecise data), model uncertainty (arising from structural and conceptual assumptions), and scenario uncertainty (linked to future system developments and normative choices). This complex interplay means that changes or gaps in one dimension often propagate through others, amplifying the overall ambiguity in outcomes.

Consequently, the coexistence of variability and diverse uncertainties complicates result interpretation and decision-making. It challenges practitioners to discern which differences are attributable to natural system fluctuations versus those emerging from data limitations or

methodological choices. Without careful acknowledgement, these overlapping factors may obscure the true robustness of findings or lead to overconfident conclusions.

To foster credible and transparent LCA practice, it is essential to explicitly recognize and communicate this complexity. This includes clearly delineating sources and types of variability and uncertainty, quantifying their impacts where possible, and discussing their implications for study conclusions. Only through such openness can LCA fulfil its role as a reliable tool for environmental decision support, guiding stakeholders with appropriate caution and contextual understanding.

2.4.6 Implications for Prospective LCA

Prospective and ex-ante LCA face unique challenges due to the need to model evolving technologies, dynamic background systems, and long-term environmental changes. These factors amplify both variability and uncertainty, complicating the interpretation and reliability of results compared to conventional LCA.

Building on foundational frameworks such as Lloyd and Ries (2007) and Huijbregts (1998a), Thonemann et al. (2020) provide a practical critique by identifying three primary methodological limitations in prospective LCA literature: comparability, data quality, and uncertainty management. Comparability suffers when model structures, functional units, or background assumptions differ without clear justification. Data quality limitations arise from reliance on poorly documented or proxy-based foreground data, especially for early-stage technologies. Uncertainty remains underexplored, particularly concerning scenario logic embedded in model choices and temporally evolving systems.

Recognizing these interconnected challenges is essential for transparent and robust prospective LCA practice. Clear communication of assumptions, explicit treatment of variability and uncertainty, and careful scenario framing help practitioners disentangle inherent system complexity from modelling decisions. This approach ultimately enhances the credibility and strategic value of prospective assessments.

The following section explores methodological advances and practical tools developed to address these challenges in prospective LCA, emphasizing their application to emerging technologies and dynamic systems.

2.5 Addressing Uncertainty & Variability in pLCA

Appropriate consideration of uncertainty and variability is essential for producing robust and meaningful prospective life cycle assessments (pLCAs). Unlike contemporary and retrospective LCAs, prospective assessments must contend with the complexities of evolving technologies, dynamic background systems, and long-term environmental changes. This section provides a structured overview of methodologies and approaches to address the key types of uncertainty—model, scenario, and parameter—as well as inherent variability across spatial, temporal, and source dimensions. By systematically confronting these challenges, practitioners can enhance the transparency, comparability, and decision-relevance of pLCA outcomes, ultimately supporting more informed sustainability strategies.

2.5.1 Addressing Model Uncertainty in LCA

Addressing model uncertainty presents distinct challenges primarily because it often lies beyond the direct control of LCA practitioners. Model structures are typically embedded within the design of impact assessment methods, databases, and software tools, which incorporate normative decisions and assumptions about environmental mechanisms and system behaviours (Huijbregts et al., 2003; Ziyadi & Al-Qadi, 2019). For example, the assumption of linearity in environmental fate models or the use of static characterization factors can introduce structural biases that are difficult to quantify and propagate using conventional uncertainty analysis techniques (Huijbregts et al., 2003; Igos et al., 2019). These assumptions lack direct empirical data or frequency-based variability, making them resistant to standard statistical uncertainty quantification methods.

Model uncertainty often reflects these non-statistical, conceptual assumptions and, as noted by Ziyadi and Al-Qadi (2019), "simple probability distributions are not capable of appropriately handling this type of uncertainty." Moreover, there is currently no universally accepted mathematical framework or standard protocol to systematically quantify model-form uncertainty in LCA. This limitation constrains the adoption of advanced methods, such as Bayesian model averaging and modular integrated modelling frameworks, which require well-defined probabilistic representations of uncertainty. These approaches are widely used in broader fields of uncertainty quantification and risk assessment to systematically incorporate multiple plausible model structures by weighting them according to evidential support (Hoeting et al., 1999).

To partially address model uncertainty, practitioners employ hybrid approaches, such as Economic Input-Output LCA (EIO-LCA). These techniques have emerged as pragmatic methods to mitigate certain aspects of model uncertainty. Baker & Lepech (2009) examined combining process-based LCA with economic input-output data to address structural uncertainties related to system boundaries and data gaps. However, this hybrid approach inherently overlooks aleatory uncertainty, further highlighting the partial nature of current mitigation strategies.

Huijbregts et al. (2003) implemented a component-specific approach, employing sensitivity analyses alongside Monte Carlo simulations and bootstrapping to explore model uncertainty in characterization factors for acidification, eutrophication, and toxic impact categories. Using alternative sets of regional and substance-specific factors, these were sampled randomly with equal probabilities across 10,000 iterations, providing a distribution of possible outcomes. This approach illustrated that ignoring model uncertainty could lead to overconfidence in LCA results. However, their quantification focused on select impact categories with available alternative data and did not address the broader range of structural uncertainties inherent in LCA modelling highlighted by Ziyadi and Al-Qadi (2019).

One notable area of advancement in addressing such structural uncertainties involves the development of dynamic characterization factors (CFs). Dynamic CFs represent a significant advancement over traditional static factors by explicitly weighting emissions according to their timing relative to environmental fate and effect processes, thereby providing a more accurate temporal representation of impacts (Levasseur et al., 2010; F. Cherubini et al., 2011). However, these dynamic CFs depend on specific model assumptions regarding pollutant transport, degradation, and bioaccumulation, which vary between approaches and lack full empirical validation, introducing new sources of model uncertainty (Hellweg et al., 2003). Parameter uncertainties, such as pollutant half-lives and effect thresholds, further contribute to variability in

impact estimates. Due to limited standardized methods and high computational demands, the robustness of temporal CFs is typically evaluated through sensitivity analyses and scenario comparisons (Igos et al., 2019).

To mitigate temporal uncertainty at the inventory level, methods such as the Enhanced Structural Path Analysis (ESPA) developed by Beloin-Saint-Pierre et al. (2014) offer a mechanistic framework that continuously tracks the timing and propagation of emissions and resource flows across complex supply chains. This approach delivers fine-grained temporal resolution within the life cycle inventory phase, reducing model and parameter uncertainty by precisely accounting for when environmental burdens occur and how they cascade upstream. Despite its methodological rigor, ESPA's practical application remains limited due to high data resolution requirements and computational demands, restricting its use primarily to research contexts.

Dynamic life cycle inventories (dynamic LCIs) originated as a more practical method to incorporate temporal variability in life cycle inventory parameters, such as emission factors and material efficiencies, by systematically updating data over the product system's lifespan based on expert projections or sectoral scenarios (Tiruta-Barna et al., 2016; Shimako et al., 2017). These methods enable a more realistic representation of evolving system states within conventional LCA databases. Dynamic LCIs enable emission factors, energy mixes, and material efficiencies to vary temporally, thereby enhancing coherence between foreground and background systems.

Building upon this foundation, recent advances have focused on integrating dynamic LCIs with scenario outputs from Integrated Assessment Models (IAMs) to provide regionally and temporally explicit prospective inventories. Tools such as THEMIS (Gibon et al., 2015), Premise (Sacchi et al., 2022), and approaches by Maes et al. (2023, 2024) automate the translation of IAM scenario data into LCI updates. While this integration improves alignment with future socio-economic pathways and energy transitions, challenges remain related to the resolution and assumptions within IAMs and the disaggregation of aggregated data into LCA-relevant process flows. Despite their advantages, dynamic LCIs remain vulnerable to model uncertainty arising from IAM scenario assumptions, data disaggregation techniques, and incomplete knowledge of emerging technologies. Hybrid modelling strategies apply dynamic LCIs selectively to key system components to balance detail with computational feasibility. It should be noted that IAMs and their internal uncertainties will be discussed in depth within Section 2.4.2, as their primary role is in shaping normative future pathways rather than structural model assumptions.

Comprehensive treatment of model-form uncertainty in LCA remains limited by methodological complexity, data scarcity, and high computational demands (Huijbregts et al., 2003; Bamber et al., 2020; Codotto et al., 2024). Popular LCA software such as SimaPro, OpenLCA, and Brightway predominantly implement uncertainty propagation for quantity (parameter) uncertainty, often using Monte Carlo simulation, while systematic quantification and propagation of model-structure uncertainty remain largely unsupported (Igos et al., 2019).

To address this challenge, practitioners rely on a set of best practices aimed at enhancing transparency and robustness despite these limitations:

- ***Systematic scenario and sensitivity analyses*** that test alternative plausible model structures or key normative choices, enabling evaluation of result robustness against structural assumptions (Huijbregts et al., 2003)

- ***Engagement with model intercomparison studies and literature syntheses*** to understand the typical magnitude and sources of model uncertainty for specific impact categories or emerging technologies (Maurice et al., 2000)
- ***Recognition of intrinsic limitations***, accepting that some model uncertainties, especially those rooted in normative decisions or gaps in scientific understanding, may be irreducible and inherently constrain the confidence of LCA results (Huijbregts et al., 2003).

Ultimately, fully quantifying model uncertainty is still an open challenge. Until more rigorous methods and software capabilities emerge, embracing transparent scenario exploration and critical appraisal of model assumptions offers the best means to enhance the reliability and clarity of LCA results in the face of structural uncertainty.

2.5.2 Addressing Scenario Uncertainty in LCA

Scenario uncertainty in LCA arises from normative choices that are unavoidable in the study design. As Huijbregts (2003) notes, these normative choices lead to uncertainty because different assumptions and boundaries can produce varying LCA outcomes. Quantifying scenario uncertainty therefore requires two steps: first, identifying and documenting the normative choices made in the LCA; second, developing methods to quantify how these choices affect output uncertainty. Building on this foundation, the following sections review key sources of scenario uncertainty relevant to prospective and consequential LCA, focusing on temporal boundaries, foreground technology development, background system dynamics, and scenario interdependencies. The discussion emphasizes contemporary methodological approaches and how they inform the present study's integration of dynamic system changes, technology scaling, and structured sensitivity analysis.

2.5.2.1 Addressing Sources of Scenario Uncertainty

Scenario uncertainty is a fundamental challenge across all forms of LCA, arising from normative choices that shape system boundaries, temporal scopes, and model assumptions. While ALCA and other conventional methods of address some aspects of scenario uncertainty, the magnitude and complexity of such uncertainty are particularly pronounced in consequential CLCA and prospective/ex-ante LCA, where the modelling of future system changes, market dynamics, and technology evolution is essential. A common approach among practitioners to assess the reliability of their models is to compare results with similar previous studies, providing valuable context and consistency checks (Igos et al., 2019). However, while this comparative method supports validation efforts, it does not fully capture the breadth of model and scenario uncertainty. Consequently, this section focuses on reviewing and critically analyzing methodologies developed to address scenario uncertainty within CLCA and pLCA frameworks, as these approaches necessitate more sophisticated and integrated methods to manage the inherent complexities and uncertainties of forward-looking assessments.

Temporal Scenario Uncertainty

Traditional LCA commonly applies a fixed time horizon, producing a static snapshot of environmental impacts based on current or projected system states. However, as Shimako et al. (2018) emphasize, "the use of a fixed time horizon does not allow the proper consideration of essential information especially for time periods encompassing the lifetime of the studied

system." This limitation is critical for products or systems with long lifespans or evolving technologies, as static boundaries obscure the timing of emissions, removals, and technology transitions that influence sustainability outcomes.

Understanding and addressing temporal scenario uncertainty, uncertainty arising from normative choices about when impacts are assessed and which future system states are considered, is essential for improving the accuracy and relevance of LCA results. Temporal uncertainty exists on multiple levels, including the timing of emissions within product systems, evolving technological states reflected in life cycle inventories (LCIs), and the selection of future time horizons for impact evaluation. These layers require distinct but complementary methodological approaches.

As discussed in Section 2.4.1, ESPA provides a detailed, mechanistic method for continuous temporal resolution within inventories but is limited in scalability. Dynamic LCIs complement this by offering a more scalable and practical approach to temporal inventory modelling through systematic updates informed by scenario data. These methods systematically update inventory data over time, incorporating projections from external scenario sources to reflect evolving system conditions. Integrated Assessment Models (IAMs) play a critical role in this process by providing coherent socio-economic, technological, and energy system pathways that inform key parameters (i.e. emission factors, energy mix compositions, and material efficiencies) over multiple future time steps.

While this foundational approach lays the groundwork for understanding temporal variability in inventories but does not itself resolve normative uncertainties about which discrete or rolling future time horizons should be used to frame impact assessments and subsequently what environmental burdens or benefits are captured. To address this higher-level scenario uncertainty, a methodological approach on the idea of a Multi-Horizon Sensitivity Analysis (MHSA) has emerged. MHSA systematically explores environmental impacts across multiple discrete or rolling future time intervals aligned with policy milestones, technology development stages, or market evolution patterns.

Although the specific phrase "Multi-Horizon Sensitivity Analysis" has seen limited use in LCA and other disciplines, the concept of assessing impacts across multiple time horizons has been recognized as essential. For example, Levasseur et al. (2010) and Beloin-Saint-Pierre et al. (2014) emphasize the need to consider multiple temporal intervals to capture the dynamic evolution of environmental impacts. More recently, Maes et al. (2021) explicitly apply this approach in pLCA, demonstrating the practical benefits of systematically evaluating environmental outcomes over several future periods. Here, we adopt the term MHSA as a concise descriptor to encapsulate these growing methodological efforts to transparently evaluate the implications of normative temporal choices on LCA outcomes.

MHSA is particularly suited to ex-ante and prospective LCA contexts where normative uncertainty centres on the selection of discrete future time horizons, specific intervals aligned with policy milestones, technology deployment phases, or market evolution patterns. It is important to note that MHSA primarily focuses on *when* impacts are assessed, while the "*what*", including marginal suppliers, technology mixes, and emission factors for each horizon, is determined by complementary modelling methods such as IAMs and marginal supplier identification approaches (discussed later in this section). MHSA acts as a time-structuring framework that defines the relevant time horizon or provides discrete "snapshots" at which the

system context must be specified and analysed. This division clarifies MHSA's role as a framework for temporal evaluation rather than a full system modelling method.

Several studies emphasize the importance of explicitly considering multiple future time horizons to capture temporal dynamics in prospective and consequential LCAs. Buyle et al. (2019) provide a comprehensive review of ex-ante LCA methodologies, underscoring the need for systematic temporal scenario analysis and sensitivity to technology learning and market adoption over time.

Maes et al. (2021) demonstrate that variations in clinker production impacts depend not only on the chosen time horizon (when), but also on evolving system conditions (what) at those times. Their ex-ante LCA shows that neglecting future technological changes, like carbon capture adoption and fossil fuel phase-out, can lead to overestimating impacts by up to five times. Early periods (2020–2030) differ significantly due to novel technology uptake. Sensitivity analysis highlights that a 20% change in electric kiln efficiency alters environmental impacts by about 10%. The study also stresses considering only plausible scenario combinations, as unlikely mixes can distort results. While careful scenario selection is crucial to maintaining result relevance, the broader implementation of MHSA introduces practical challenges that must be managed.

MHSA offers substantial advantages in elucidating temporal uncertainties; however, it is computationally demanding and reliant on temporally resolved, reliable data that may be scarce for emerging technologies. Moreover, presenting a spectrum of temporal outcomes can complicate stakeholder communication, necessitating careful framing to distinguish temporal variability from mere uncertainty (Buyle, Audenaert, et al., 2019). Despite these challenges, MHSA remains a pragmatic and widely adopted tool for enhancing the transparency and interpretability of temporally sensitive LCA outcomes, particularly in ex-ante decision-making contexts where normative temporal choices critically influence sustainability assessments. The systematic generation of coherent, temporally explicit scenarios required for MHSA is increasingly facilitated by Integrated Assessment Models (IAMs), which provide harmonized socio-economic, technological, and environmental pathways.

Scenario Consistency and Integrated Assessment Models (IAMs)

Integrated Assessment Models (IAMs) have become foundational in advancing scenario consistency and reducing uncertainty in prospective and consequential LCAs. These models synthesize climate science, economics, energy system transitions, land use, and policy frameworks into comprehensive, internally consistent scenario narratives spanning multiple decades and geographic scales.

IAMs as Scenario Engines for Prospective LCAs

pLCAs depend heavily on the quality and consistency of scenario inputs to project future system states accurately. IAMs, including prominent frameworks such as IMAGE, REMIND-MAgPIE, and MESSAGE-GLOBIOM, produce harmonized projections of greenhouse gas emissions, energy supply mixes, technology deployment, and policy impacts aligned with global climate targets (Van Vuuren et al., 2011; Harmsen et al., 2021; Luderer et al., 2015; Maes et al., 2023). It is important to distinguish the role of IAMs in defining what system components and processes exist at specific future points in time, the aforementioned 'what' and 'when', respectively, from

the subsequent translation of these components' environmental loads into impact metrics, the 'how'.

IAM outputs define the LCI by specifying future energy mixes, technology adoption, material production, and other system parameters at specific time points. Primarily, IAMs modify background system inventories, such as electricity generation, fuel production, and upstream material supply chains, reflecting structural changes in these sectors over time. Foreground system components, representing specific processes under study, are generally defined separately and may require independent modelling or adaptation.

This information shapes both foreground and background inventories in pLCAs. In contrast, the LCIA phase translates these inventory flows into environmental impacts using characterization factors and impact metrics. LCIA primarily involves methodological modelling choices that are separate from IAM scenario data. While IAMs may indirectly influence LCIA assumptions through expected policy or scientific trends, the calculation of impacts and their temporal weighting remain distinct model decisions beyond the scope of IAM outputs.

Maes et al. (2023, 2024) provide a salient example of this integration, systematically incorporating IAM-derived scenario data into prospective LCAs of the cement and electricity sectors. Their work goes beyond technical data integration, emphasizing the critical role of IAMs in ensuring scenario consistency by employing harmonized, predefined socio-economic and climate pathways. These internally coherent scenario narratives enhance replicability and reduce uncertainty by grounding LCAs within structured, policy-relevant backbones, thereby avoiding contradictory assumptions across sectors and timeframes.

This consistency is essential for producing credible and actionable sustainability assessments aligned with global climate targets. While integrating IAM data poses challenges, such as spatial and sectoral resolution mismatches and the complexity of disaggregating aggregated outputs, Maes et al. demonstrate that careful methodological approaches can mitigate these issues effectively. Notably, their 2023 study quantifies substantial benefits of IAM integration, showing that incorporating evolving marginal electricity mixes from IAM scenarios results in greenhouse gas emission reductions of 50–75% over long-term horizons, reflecting anticipated grid decarbonization pathways. These findings highlight how scenario-consistent inventories improve the fidelity and temporal sensitivity of pLCA outcomes, supporting more robust decision-making under uncertainty. Beyond global warming potential, the approach offers a framework for assessing the dynamic interplay of technology adoption, policy shifts, and market evolution, representing a significant advancement in the application of IAMs within environmental assessment frameworks.

Similarly, Sacchi et al. (2022) developed the Premise framework, an advanced tool that automates the translation of IAM scenario outputs into regionally and temporally explicit dynamic LCIs. Premise integrates hybrid process-based and environmentally extended input-output (EEIO) LCA models, enabling comprehensive and scalable prospective assessments across complex industrial and energy sectors. The framework employs econometric and mapping algorithms to disaggregate aggregated IAM outputs, often presented at coarse spatial and sectoral scales, into detailed LCI flows that reflect regional technology mixes, energy carriers, and production processes over multiple future time steps.

In pCLCA it is defensible, and often preferable, to separate (i) the construction of a frozen, scenario-consistent background using IAM-to-LCI translation (e.g., premise) from (ii) the identification of future marginal suppliers for displacement. The former yields internally coherent inventories for a chosen year and pathway (Sacchi et al., 2022); the latter can be implemented transparently in the foreground when the credited product market requires finer-grained definitions than the IAM-to-LCI pipeline provides. This division of labour is consistent with recent guidance formalizing marginal-supplier identification from IAM trajectories (choice of window, foresight/lead times, capital replacement, and growth measurement) while leaving the market definition to the practitioner (Maes et al., 2023).

The two steps can also be combined when a sector is adequately handled by the IAM-to-LCI pipeline. Specifically, where premise's consequential parameterization for a sector (e.g., electricity or fuels) yields a marginal supplier mix that matches the credited product's market definition (functional equivalence, grade/spec, geography) and its settings (time window, foresight/lead time, lifetime/capital replacement, growth metric) are disclosed, the premise-derived marginal mix may be adopted directly. For product-grade markets not resolved by IAM mappings (e.g., extrusion-grade aluminium billet, sub-national routing), marginal identification remains in the foreground, implemented transparently with evidence-tier rules and trade-routed exporter bins. This hybrid approach preserves a frozen, scenario-consistent background while aligning marginal supplier identification with recent guidance on using IAM trajectories.

Presently, premise has been applied across diverse sectors, including cement, steel, transportation, and power generation, where it improves the temporal and spatial resolution of inventory data and better represents emerging technologies and decarbonization pathways. This hybrid modelling approach captures the intricate interdependencies within supply chains while maintaining computational feasibility, a critical balance for prospective LCA studies.

Sacchi et al. acknowledge significant challenges in harmonizing IAM data with LCA inventories, including the need to reconcile sectoral aggregation mismatches, regionally varying technology adoption rates, and temporal consistency across datasets. Premise addresses these issues through iterative data disaggregation, alignment with industry roadmaps, and expert input to refine scenario translation fidelity. Additionally, the framework incorporates dynamic updates to inventories that reflect evolving market conditions and regulatory landscapes, facilitating continuous scenario integration. The integration of IAM scenarios into LCAs follows a multi-step process that is described in Sacchi et al. (2022) and depicted in Figure 2-6.

IAM-driven prospective LCAs offer enhanced scenario consistency and improved representation of evolving system dynamics, yet they also introduce specific challenges. Mismatched resolution occurs because IAM outputs are generally designed for global or regional analyses with coarse sectoral aggregation, while LCAs often require fine-grained, process-level data. Bridging this gap requires careful data disaggregation, which can introduce uncertainties and methodological biases (Sacchi et al., 2022).

Furthermore, incorporating time-series data for multiple sectors and regions increases data management complexity and computational demands, sometimes limiting practical applicability. They are also subject to scenario dependence, embodying assumptions about technology costs, policy effectiveness, and socio-economic trends. Divergence among IAMs and scenario families necessitates multi-scenario exploration to bracket uncertainty (Riahi et al., 2011). Despite these

challenges, the consensus in recent literature underscores IAM integration as a best practice for managing scenario uncertainty in prospective LCAs.

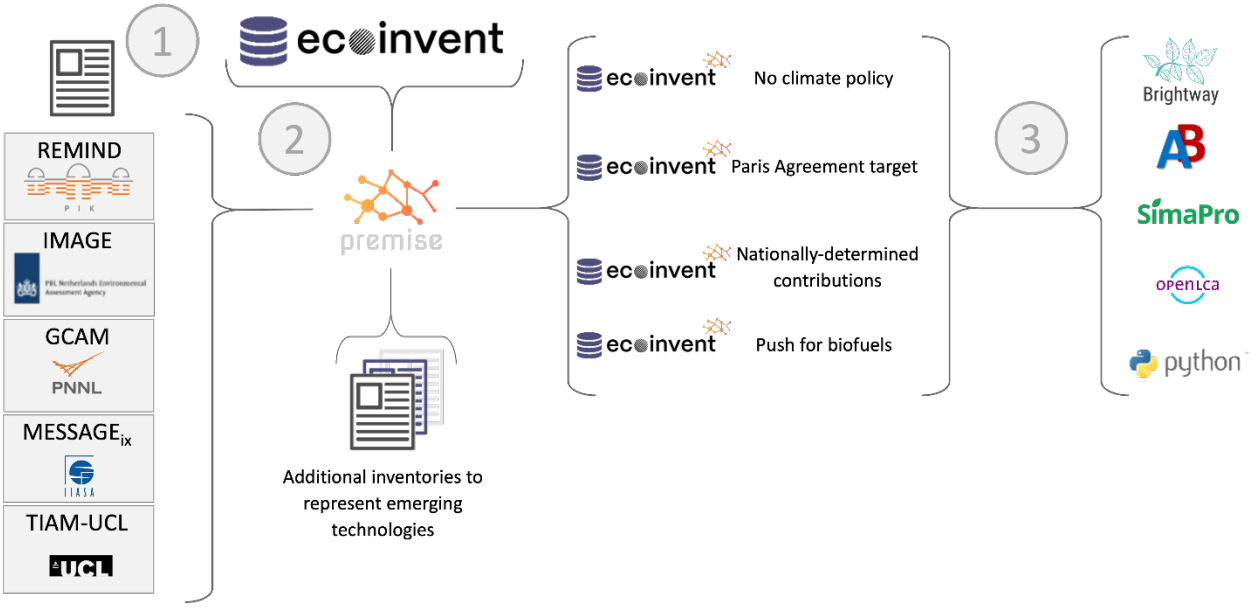


***Figure 2-6**: Schematic of the IAM–LCA coupling workflow used to generate prospective background databases derived from (Sacchi et al., 2022). (1) Integrated assessment model outputs (REMIND, IMAGE, GCAM, MESSAGEix, TIAM-UCL) provide scenario trajectories that (2) are applied to the baseline ecoinvent database via premise to create policy- and scenario-specific inventories (e.g., no climate policy, Paris Agreement target, nationally determined contributions, biofuel push), optionally augmented with additional inventories for emerging technologies. (3) The resulting scenario-specific databases are then used in LCA tools (e.g., Brightway/Activity Browser, SimaPro, openLCA, Python).*

Marginal Supplier Identification and Market Delimitation

Marginal supplier identification is a critical component in consequential life cycle assessment (CLCA). It determines which producers or technologies respond to changes in demand, thereby influencing the resulting environmental impact calculations. This process typically involves delineating market boundaries, both geographical and sectoral, to define the relevant supply context, and analyzing economic or market factors such as price elasticity, spare capacity, and cost competitiveness to identify suppliers likely to adjust production in response to demand shifts. In contemporary or current CLCA studies, marginal suppliers are typically identified based on static or near-term market conditions, assuming relatively stable technology portfolios and limited temporal variation (refer to Section 2.0). This approach supports impact assessments grounded in the present-day supply chain landscape (B. Weidema, 2003; Ekvall & Weidema, 2004).

To navigate the complexity and scale of modern supply systems, particularly when detailed data are sparse or systems are large, heuristic methods have become a practical cornerstone for marginal supplier identification within CLCA (Earles & Halog, 2011; Maes et al., 2023). Notably, Weidema's four-step heuristic procedure provides a systematic and transparent approach to narrowing the pool of potential marginal suppliers (B. Weidema, 2010). This method

involves delineating relevant market boundaries, assessing supplier capacities, evaluating production costs, and considering market share dynamics to determine the most likely marginal producers (see Figure 2-7). Its scalability and adaptability have made it widely used across diverse sectors and applications.

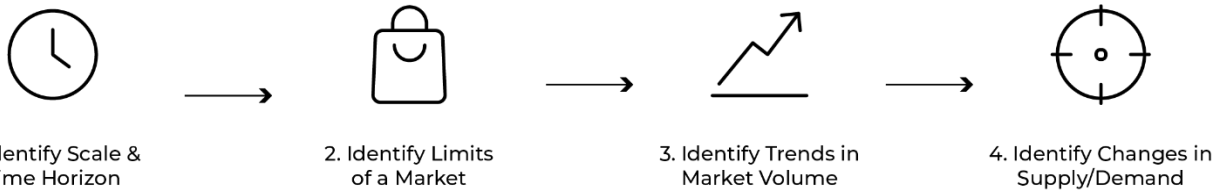


***Figure 2-7**: Four-step heuristic for identifying long-term marginal supplier responses in consequential LCA (derived from (B. Weidema, 2010))*

Building upon this, ecoinvent database's consequential modelling framework exemplifies a pragmatic hybrid approach, combining heuristic decision rules (i.e. market delimitation, capacity assessment, and cost ranking) with empirical market data including observed dispatch orders, capacity factors, and regional energy mixes (ecoinvent Association, 2025). This integration allows the model to leverage systematic, rule-based supplier identification while grounding assumptions in actual market behaviour and regional specifics, enhancing both scalability and contextual accuracy.

In the case of electricity datasets, ecoinvent determines marginal mixes by considering capacity utilization, production costs, and typical dispatch orders, thereby aligning heuristic principles with observed market behaviours (Vandepaer et al., 2019). This method provides region-specific marginal supplier mixes reflecting current market conditions, though it largely assumes static market structures, limiting its applicability in rapidly evolving contexts.

This approach provides region-specific marginal supplier mixes to a meaningful degree by incorporating empirical market data, dispatch patterns, and observed operational behaviours where available. However, for many non-electricity processes, data limitations and granularity constraints often necessitate the extrapolation of information from regions with better coverage—such as applying European market characteristics to North American contexts—which can introduce regional biases (Wernet et al., 2016). While such manual refinements enhance relevance beyond generic or global datasets, these methods generally assume relatively static market structures and supply conditions. Consequently, they may not fully capture rapid technological advancements, capacity expansions, or evolving market dynamics, limiting their applicability in sectors experiencing significant transformation or emerging markets.

However, as CLCA methodologies evolve towards prospective CLCA (pCLCA), the need arises to explicitly capture temporal, spatial, and technological dynamics that characterize future supply systems. Unlike traditional CLCA, pCLCA often needs to consider the evolution of markets over extended periods of time, including technological innovation, capacity expansions, policy interventions, and shifting economic contexts (Sacchi et al., 2022; Maes et al., 2024; Nurdiawati et al., 2025). These dynamics necessitate the adaptation of marginal supplier identification methods to reflect not only the current supply configuration but also plausible future states. To ensure robustness, this adaptation is increasingly informed by IAMs that drive dynamic market and supplier modelling (Maes et al., 2024).

To effectively capture future market evolutions, recent research advocates extending heuristic methods by integrating industry roadmaps, technology adoption curves, and expert knowledge to project the dynamic composition of marginal suppliers (Hertwich et al., 2015; Georgiades et al., 2023; Maes et al., 2023). These projections allow the heuristic framework to move beyond static snapshots, incorporating anticipated capacity expansions, technological improvements, and policy-driven transitions into marginal supplier identification. By grounding heuristic rules in structured forward-looking data, such as industry roadmaps and integrated assessment scenarios, prospective consequential LCA can more accurately reflect plausible future supply configurations, thereby improving the relevance and robustness of impact assessments.

A compelling example of prospectively utilizing the heuristic method's adaptability is its application to the Danish electricity market, where the approach is tailored to reflect the region's unique temporal and spatial grid characteristics. Muñoz and Weidema (2023) demonstrate how this tailored heuristic approach captures the marginal electricity mix over multiple years, reflecting shifts in generation technologies and dispatch order. Quantitatively, their analysis shows the marginal supplier composition evolving from predominantly fossil-fuel-based units to increased shares of renewables and flexible generation, resulting in a projected reduction of marginal greenhouse gas emissions by approximately 40% over a decade. This flexibility illustrates how the heuristic framework balances methodological rigor with contextual specificity, supporting its broad applicability across diverse sectors and regions.

Recognizing the complexity of real-world markets, later advances incorporated economic modelling techniques—particularly input-output (IO) and hybrid LCA approaches—that improved the representation of upstream supply chains and indirect effects by leveraging economic transaction data. These methods, reviewed extensively by Weidema et al. (2018) and Mendoza Beltrán et al. (2020), enhanced the granularity and completeness of life cycle inventories. However, they generally maintain static assumptions about market relationships and do not explicitly simulate marginal supplier shifts in response to incremental demand changes, which limits their ability to fully capture consequential dynamics. More complex economic equilibrium models, including partial and general equilibrium frameworks, offer theoretical rigor by modelling market price and supply adjustments but remain computationally intensive and are less frequently applied in practice due to data and resource constraints.

However, accurately identifying marginal suppliers and defining market boundaries in pCLCA requires addressing scenario uncertainty—the uncertainty about which future socio-economic, technological, and policy pathways will unfold. This involves selecting plausible scenarios that represent alternative futures and acknowledging that different scenarios lead to fundamentally different marginal supplier compositions and market evolutions. Once scenarios are selected, the primary sources of uncertainty shift to parameter uncertainty, which arises from variability in inputs within those scenarios such as technology adoption rates, capacity expansions, and policy stringency. While parameter uncertainty and its mitigation strategies are discussed in detail in Section 2.4.3, it is important to recognize the inherent overlaps between the structural components of LCAs and the uncertainties they entail.

Scenario uncertainty is mitigated by exploring multiple distinct scenarios, each embodying different assumptions about technological progress, market dynamics, and policy interventions. Sensitivity testing of these alternative scenarios allows analysts to bracket the range of possible outcomes and understand the implications of different future pathways. A key dimension of

scenario uncertainty involves defining spatial and sectoral market boundaries. Selecting appropriate geographic scopes and relevant competing technologies often depends on expert judgment and extrapolation from incomplete or regionally biased data, particularly in emerging markets (Wernet et al., 2016). Sensitivity analyses around these market delimitation choices provide valuable understanding of how marginal supplier identification responds to different assumptions, enhancing confidence in the robustness of results.

Mendoza Beltrán et al. (2020) illustrate that, while individual scenarios can be analysed in detail, capturing the full scope of scenario uncertainty necessitates comparing multiple well-defined futures. This highlights the importance of transparent and deliberate scenario selection combined with scenario-based sensitivity testing as essential strategies to manage uncertainty at the system level.

Furthermore, integrating foreground marginal supplier models with IAM projections—such as through the Premise framework (Sacchi et al., 2022)—facilitates scenario consistency by harmonizing assumptions about background system transformations and socio-economic pathways. Despite this harmonization, IAMs inherently represent multiple potential futures, and the uncertainty regarding which pathway will prevail must be explicitly acknowledged and managed through multi-scenario analyses.

Together, these practices reflect a progression from static, single-scenario heuristics toward multi-scenario, scenario-explicit frameworks that embrace uncertainty as a core consideration. Such an approach improves the credibility and policy relevance of marginal supplier identification in pCLCA, supporting more resilient and informed decision-making under complex, uncertain future conditions as failing to account for these temporal and market dynamics risks misattributing environmental burdens or benefits, particularly when novel technologies are introduced and must compete with incumbent processes (Maes et al., 2021).

Scenario Interdependencies and Embedded Normativity

Scenario uncertainty in prospective and consequential Life Cycle Assessment (pLCA and CLCA) is rarely the product of isolated assumptions. Rather, it arises from a complex web of interdependent choices and normative frameworks that collectively shape the construction and interpretation of future scenarios. Recognizing and explicitly modelling these interdependencies is crucial to improving transparency, robustness, and validity in environmental impact assessments.

Interdependencies occur when assumptions regarding one aspect of the scenario inherently influence or constrain others. For example, the choice of time horizon affects which technologies are considered mature or emerging in the foreground system, which in turn determines the composition and dynamics of the background system such as energy mixes or material supply chains. Similarly, socio-economic pathways derived from IAMs impose constraints and feedback on technological deployment rates, policy implementation, and market behaviour. This systemic interconnectedness means that scenario elements cannot be treated as independent variables but must be jointly considered within an integrated modelling framework (Bamber et al., 2020)

This notion of structural scenario interdependencies has been extensively discussed in the literature. Bamber et al. (2020) highlight how seemingly discrete scenario choices cascade to shape trajectories of modelled futures, necessitating causal reasoning to understand their

compounded effects. Mendoza Beltran et al. (2020) emphasize that the interwoven nature of assumptions in pLCA requires scenario designs that embed normative expectations within logical and temporal consistency constraints. These perspectives underscore the importance of scenario coherence to avoid contradictory or unrealistic futures that could bias LCA outcomes.

Methodologically, addressing interdependencies involves employing integrated scenario construction frameworks that explicitly encode causal relationships and feedback loops among scenario components. Techniques such as causal loop diagrams, system dynamics modelling, or scenario logic frameworks (e.g., the SIMPL framework) have been applied to maintain internal consistency and facilitate transparent documentation of normative choices (Maes et al., 2024). Furthermore, scenario-based sampling methods allow systematic exploration of correlated assumptions, enabling sensitivity and uncertainty analyses that respect the structural dependencies inherent in scenario narratives.

Despite methodological advances, challenges remain. Modelling interdependencies increases computational complexity and data requirements. Moreover, the embedding of normative choices (i.e. values, preferences, or policy priorities) into scenario logic can introduce subjectivity that complicates interpretation and communication. Transparency in documenting these choices and their interrelations is therefore essential to preserve credibility and reproducibility.

2.5.3 Addressing Parameter Uncertainty in LCA

Building upon the previously identified sources of parameter uncertainty this section focuses on the methodological approaches used to manage and reduce these uncertainties in pCLCA. Parameter uncertainty is not unique to pCLCA; it has long been recognized as a fundamental challenge across all LCA types. However, in pCLCA, these uncertainties are often exacerbated by evolving technologies, dynamic supply chains, and shifting policy contexts. Despite this increased complexity, key foundational methods like data quality indicators, pedigree matrices, probabilistic modelling, Monte Carlo simulation, and sensitivity analysis have been established since early seminal works (Weidema & Wesnæs, 1996; Huijbregts et al., 2003; Lloyd & Ries, 2007) and continue to underpin contemporary studies, including pCLCA.

Given the complexity and volume of parameters involved in typical LCA models—often numbering in the thousands—it is impractical to characterize uncertainty for every input in detail. Instead, effective uncertainty treatment requires systematic methods to identify, quantify, and prioritize parameters that most significantly affect model outcomes.

This section outlines key approaches employed to treat parameter uncertainty in prospective and consequential LCAs. It discusses strategies for assessing data quality, applying probabilistic sampling techniques, performing sensitivity analyses, addressing database and structural limitations, and accounting for spatial and temporal variability. The goal is ultimately to enhance transparency and reduce uncertainty in LCA results, thereby supporting more informed decision-making.

2.5.3.1 Addressing Sources of Parameter Uncertainty

Probabilistic Modelling and Monte Carlo Simulation

Probabilistic modelling is a foundational method for managing parameter uncertainty in LCA, particularly within pCLCA where temporal and technological dynamics introduce significant

variability. Unlike traditional deterministic approaches that assign single fixed values to input parameters (Lloyd & Ries, 2007), probabilistic modelling represents uncertain inputs as probability distributions—such as normal, lognormal, or triangular—reflecting observed data variability, expert judgment, or literature-derived ranges (Cox et al., 2020). This probabilistic representation captures both inherent variability and epistemic uncertainty in LCI data and impact assessment factors, providing a more realistic characterization of uncertainty.

The prevalence and established nature of Monte Carlo simulation (MCS) in LCA uncertainty analysis is well documented (Groen et al., 2014a; Bamber et al., 2020). For example, Lloyd and Ries (2007) reviewed 16 stochastic LCA studies and found that 10 explicitly used MC with random sampling. Other sampling techniques included Latin Hypercube Sampling, parametric bootstrapping, and combinations of multiple approaches. While justifications for choosing specific sampling methods were not always discussed, MC with random sampling remains the dominant choice due to its straightforward implementation and robustness.

The appropriate number of sampling iterations required to achieve reliable results has been widely studied and varies considerably in practice. Lloyd and Ries (2007) reported iteration counts in LCA studies range from as few as 100 (Kennedy et al., 1996) to upwards of 30,000, with an average of approximately 7,000 iterations. This variation reflects an evolving understanding of the trade-off between computational cost and statistical convergence, where iteration counts directly influence the stability and reproducibility of uncertainty estimates. Over time, practitioners have recognized the increasing need for larger sample sizes—often in the range of 1,000 to several thousand iterations—to ensure robust estimation of variances and tail behaviours critical for uncertainty characterization. The computational intensity of Monte Carlo simulation, especially at high iteration counts and with large numbers of input parameters, has motivated the development of alternative and complementary uncertainty propagation methods that can improve efficiency without sacrificing robustness. Among these, Latin Hypercube Sampling (LHS) and quasi-Monte Carlo (QMC) sampling have gained prominence. LHS stratifies the input space to ensure more uniform coverage with fewer samples, thereby reducing the number of required iterations for convergence compared to simple random sampling (Lloyd & Ries, 2007; Groen et al., 2014b). Similarly, QMC uses low-discrepancy sequences to systematically fill the input space, improving convergence rates and decreasing computational burden (Groen et al., 2014b; Bamber et al., 2020).

Beyond sampling-based methods, analytical uncertainty propagation techniques offer an alternative by mathematically deriving output uncertainty based on input distributions and model functions, thereby circumventing the need for large numbers of simulations (Lloyd & Ries, 2007). While computationally efficient, these approaches require tractable model formulations and may struggle with highly nonlinear or complex LCA models. Another less commonly applied but conceptually important method is fuzzy interval arithmetic, which models uncertainty with ranges and membership functions rather than probability distributions. This approach can handle imprecise or subjective uncertainty but has limited adoption in LCA due to interpretability challenges and integration complexity (Bamber et al., 2020).

Collectively, these alternative methods complement traditional MC simulation by offering options that balance computational feasibility, data availability, and model complexity. Selection among them depends on the specific LCA context, data quality, and desired output fidelity. Incorporating these alternatives alongside MC enables analysts to tailor uncertainty propagation

to their study's unique needs while managing computational resources effectively. Within pCLCA, MC simulation is particularly valuable due to the complex and dynamic nature of the systems assessed.

Data Quality Assessment and Pedigree Matrices

Data quality fundamentally determines the credibility and usefulness of LCA results. Poor or uncertain data can obscure true environmental impacts and mislead decision-making. Therefore, evaluating data reliability, completeness, and contextual relevance is essential to properly manage parameter uncertainty. Data Quality Indicators (DQIs) form the foundational framework for evaluating the quality of inventory data in LCA. They represent qualitative or semi-quantitative criteria used to assess critical aspects of data quality such as reliability, completeness, temporal correlation, geographical representativeness, technological relevance, and sample size. DQIs help analysts systematically describe the strengths and weaknesses of data sources, providing essential context for interpreting uncertainty and guiding data improvement efforts (Fava et al., 1994; Singhofen et al., 1996).

Building upon this conceptual framework, the Pedigree Matrix was developed by Weidema & Wesnæs (1996) as a practical tool to operationalize DQIs for quantitative uncertainty analysis. The pedigree matrix translates expert judgments about each data quality criterion into ordinal scores, typically ranging from 1 (high quality) to 5 (low quality) across the following aspects: reliability of data sources, reflecting their trustworthiness and precision; completeness, which considers the coverage of all relevant flows and processes; temporal correlation, or how well the data aligns with the study's time frame; geographical correlation, indicating the representativeness of data for the spatial context; and technological correlation, which evaluates the applicability of data to the specific technologies or processes being modelled. This means that instead of treating all input data as perfectly accurate fixed values, the pedigree matrix adjusts the confidence placed in each data point by effectively "inflating" or "deflating" its uncertainty range according to its assessed quality. For example, data with high reliability and completeness receive lower uncertainty factors, indicating higher confidence, whereas data with poor documentation or limited relevance receive higher uncertainty factors, signalling greater potential variability or error.

When incorporated into uncertainty propagation methods like Monte Carlo simulation, these adjusted uncertainty factors influence the sampling distributions of input parameters, leading to more realistic and nuanced modelling of the overall uncertainty in the LCA results. This process helps avoid overconfidence in low-quality data and ensures that uncertainty estimates reflect the true quality and reliability of the underlying information. This approach has been widely integrated into LCA databases and software, enabling transparent uncertainty reporting tied to data quality (Lloyd & Ries, 2007). Several studies have successfully implemented the pedigree matrix to enhance uncertainty analysis in diverse contexts; for example, Georgiades et al. (2023) applied the approach to quantify and propagate uncertainty related to clinker production within the cement manufacturing process, demonstrating how pedigree-based assessments can improve the transparency and robustness of environmental impact estimates in complex industrial systems.

While effective, pedigree matrix methods depend on expert judgment, which may introduce variability and bias. Additionally, standardized scoring may inadequately reflect sector-specific or regional nuances, potentially misestimating uncertainty—for example, "high-quality" scores

may mean different things across sectors with differing data variability (Bicalho et al., 2017; Bamber et al., 2020). This approach has been further refined and adopted in various studies and integrated into prominent LCA databases such as ecoinvent (Kennedy et al., 1996; Frischknecht et al., 2007). For example, Guo and Murphy (2012) explain how the pedigree matrix transforms data quality indicators into 'default' lognormal probability distributions, allowing transparent quantification of data uncertainty linked directly to data quality assessments.

This evolution is particularly important in prospective and emerging technology LCAs, where empirical data are often scarce. In these contexts, hybrid data quality frameworks provide a more rigorous and transparent way to communicate data limitations and guide targeted research priorities. When incorporated within probabilistic uncertainty frameworks, such integrated assessments significantly enhance the credibility and interpretability of LCA results, especially when combined with sensitivity and scenario analyses (Ravikumar et al., 2018).

Sensitivity Analysis

Sensitivity analysis complements probabilistic modelling by systematically evaluating how variations in input parameters affect LCA model outputs. This process helps identify which parameters most strongly influence environmental impact results, guiding targeted data collection, model refinement, and uncertainty reduction efforts, thereby enhancing study robustness and efficiency (Ravikumar et al., 2018; Igos et al., 2019).

Sensitivity analysis encompasses a range of methods to evaluate how variations in input parameters affect model outputs. Local sensitivity analysis (LSA), commonly referred to as a "one-factor-at-a-time" (OAT) approach, varies individual inputs within plausible ranges while holding others constant. This method provides straightforward insights into the marginal influence of each parameter but may fail to capture nonlinear interactions and combined effects that are common in complex environmental models (Lacirignola, 2017). In contrast, global sensitivity analysis (GSA) evaluates parameter influences over their entire uncertainty ranges and explicitly accounts for interactions and nonlinear responses. Techniques such as Sobol indices, variance-based decomposition, and squared Spearman rank correlation coefficients quantify each parameter's contribution to output variance, enabling a comprehensive prioritization of key drivers (Groen et al., 2014a).

For example, Groen et al. (2014a) demonstrated that Sobol' indices provide robust sensitivity estimates across both small and large uncertainty contexts and offer more detailed insights than simpler approaches such as key issue analysis. However, full variance-based GSA can be computationally expensive, especially for models with many parameters. To address this, the Method of Elementary Effects (MEE), also known as the Morris method (MD, 1991), is widely used as a computationally efficient screening tool in LCA. MEE assesses the influence of parameters by sampling multiple points within their ranges and calculating 'elementary effects'—the change in output due to parameter perturbation—while capturing potential nonlinearities and interactions through the variability of these effects (Groen et al., 2014a; Igos et al., 2019).

Mutel et al. (2013) provide a clear example of the Morris method's practical value, successfully screening a large number of input parameters to identify those most influential on model outcomes, allowing focused refinement efforts. Despite these advances, local OAT methods remain common due to their simplicity and ease of interpretation, though they risk

underestimating parameter importance when nonlinearities or interactions are present (Groen et al., 2014a). Sensitivity analysis is particularly valuable in prospective and emerging technology assessments, where input data often have high variability or limited empirical validation. By revealing which assumptions—such as technology adoption rates, process efficiencies, or policy impacts—exert the most influence on environmental results, sensitivity analysis informs focused research and data collection to reduce uncertainty in future iterations (Hetherington et al., 2014).

Hetherington et al. (2014) advocate for the integration of sensitivity analysis in LCAs of emerging technologies to address data and system uncertainties, although their study does not include a formal sensitivity analysis. In contrast, Arvidsson et al. (2015), conducted a detailed sensitivity analysis on cellulose nanofibril production routes, identifying electricity consumption during mechanical treatment and solvent recovery efficiency as the dominant parameters affecting global warming potential (GWP). These results underscore the significant role of energy-intensive processes in driving environmental impacts.

Hybrid and Expert Elicitation Approaches

Hybrid approaches combine empirical data, literature review, and expert judgment to address parameter uncertainty, particularly when direct measurements or comprehensive datasets are lacking. These methods are especially valuable for emerging technologies or novel processes where data scarcity and variability present significant challenges (Lacirignola, 2017; Arvidsson et al., 2018).

Expert elicitation involves structured consultation with domain specialists to estimate uncertain parameters, validate assumptions, and identify critical knowledge gaps. Techniques such as Delphi surveys, structured interviews, and consensus workshops facilitate the aggregation of expert opinions, allowing for the development of probabilistic parameter distributions grounded in practical experience rather than limited data alone.

By integrating expert input with available quantitative data, hybrid approaches improve parameter estimation while capturing epistemic uncertainties arising from incomplete knowledge. For example, in scaling laboratory results to industrial processes, experts provide insights on process variability, technological learning rates, and operational constraints that may not yet be documented in the literature (Arvidsson et al., 2023). Similarly, in prospective LCAs of bio-based systems or advanced materials, hybrid methods help define plausible parameter ranges and scenario envelopes that reflect both technical potential and market realities (Thonemann et al., 2020).

Embedding expert elicitation within probabilistic frameworks strengthens uncertainty analyses by combining qualitative knowledge with quantitative modelling, offering a more comprehensive representation of parameter uncertainty (Cucurachi et al., 2018). This approach also facilitates iterative model refinement as new data emerge, and expert consensus evolves.

2.5.4 Addressing Variability in LCA

Variability represents the inherent, real-world heterogeneity in environmental conditions, technologies, and operational practices across spatial, temporal, and source-based dimensions (Heijungs & Huijbregts, 2004; Marsh et al., 2023). Unlike uncertainty, which arises from incomplete knowledge and can be propagated through probabilistic modelling, variability is

intrinsic to the system and cannot be reduced or propagated directly within LCAs (Huijbregts, 1998a).

Despite this, variability is acknowledged and indirectly incorporated in LCA through stratification and descriptive analyses. Specifically:

- ***Spatial Variability*** which reflects real-world geographic differences in environmental sensitivity and emissions, is approximated in LCA through regionalized LCIA methods and spatially explicit inventories. These methods apply location-specific characterization factors or aggregate data by broad environmental regions, capturing major spatial differences while smoothing out finer-scale heterogeneity due to data limitations and computational constraints (C. Mutel et al., 2019).
- ***Temporal Variability***, encompassing fluctuations and long-term trends in system operations and emissions, is represented by selecting appropriate temporal resolutions and employing archetypal days or representative time slices. This approach balances the need to capture key temporal dynamics with computational feasibility, providing an efficient approximation of temporal patterns without modelling continuous time series in full detail (Amor et al., 2014; Lueddeckens et al., 2020).
- ***Source-Based Variability***, arising from heterogeneity among data sources, technologies, or operational conditions, is managed by compiling multiple datasets to represent the plausible range of process performances. Sensitivity and uncertainty analyses over these datasets further characterize variability, enabling LCA practitioners to approximate real-world operational differences without requiring exhaustive data on every individual source (Cucurachi et al., 2018).

Variability interacts with parameter, model, and scenario uncertainties, complicating result interpretation and decision-making. Consequently, LCA practitioners emphasize transparent communication that clearly distinguishes variability from uncertainty, often incorporating variability considerations into uncertainty propagation through model or scenario-based approaches (Huijbregts, 1998a).

2.6 Summary & Research Gap

The literature points to three intertwined dimensions of LCA, causality, technological maturity, and time, that must be handled explicitly for forward-looking assessments. CLCA captures system responses to decisions, but without temporal framing it remains a static snapshot. Prospective approaches address this by aligning inventories with evolving backgrounds (e.g., power mixes, material markets) and by structuring futures through scenarios. Across this landscape, uncertainty appears in three forms, model, scenario, and parameter, with scenario uncertainty (normative choices about futures) and background alignment emerging as the dominant challenges. IAMs enable internally consistent futures and support dynamic LCIs, while methods such as discrete multi-horizon analysis help expose sensitivity to the choice of time horizon. Marginal supplier identification remains pivotal for CLCA yet is often handled with static heuristics that under-represent future market shifts.

Despite progress, the built-environment literature still lacks applications that (i) couple prospective CLCA with IAM-aligned, time-evolving LCIs, (ii) treat end-of-life pathways and secondary material markets as first-order drivers, (iii) ensure scenario consistency between

foreground and background systems, and (iv) separate scenario from parameter uncertainty in a transparent way—particularly in a Canadian context. This thesis addresses the gap by:

- Developing a pCLCA framework that aligns foreground modelling with IMAGE-based background projections mapped to ecoinvent.
- Focusing on end-of-life dynamics and emerging recycling options for a curtain wall system, including clear system-expansion logic.
- Applying consequential linking for marginal suppliers with targeted, documented overrides at identified hotspots.
- Using a discrete multi-horizon analysis with a near-term reference case and a longer-term stringency sweep across mitigation pathways.
- Propagating parameter uncertainty probabilistically and reporting scenario spread separately for clarity and decision relevance.

Together, these choices operationalize lessons from the literature into a tractable, policy-relevant approach that improves temporal fidelity, causal consistency, and uncertainty transparency for prospective assessments in the built environment.

Chapter 3

3 Methods

3.1 Methodological Framework Overview

The following sections introduce a novel methodological framework that advances prospective consequential LCA (pCLCA). The framework built on recent IAM–LCA coupling approaches (Sacchi et al., 2022; Maes et al., 2023) by integrating evolving background systems with detailed, probabilistic modelling of both contemporary and emerging end-of-life (EoL) technologies—an area underrepresented in current literature (Mendoza Beltran et al., 2020). The approach leveraged standardized databases, open-source tools, and harmonized assumptions to promote transparency and reproducibility. The framework treated foreground and background dynamics as interlinked, ensuring that technological, market, and environmental changes were captured consistently across scenarios.

To improve readability and reproducibility, the methods and application of this study were organized as rules first, choices later. Chapter 3 (Methods) set out these rules (i.e., universal accounting policies, data-transformation procedures, and modelling conventions that apply to any application), whereas Chapter 4 (Case Study) described choices related to the concrete implementation (e.g., specific functional unit, geography, booking years, scenario set, hotspot markets, and figure set). The analysis first constructed scenario backgrounds, then derived marginal suppliers from within those backgrounds, and finally modelled EoL pathways and their Stage D substitutions. Uncertainty was treated explicitly as two complementary layers: scenario uncertainty from alternative futures and parameter uncertainty from both background inventories and foreground EoL models. This was implemented through two distinct but related workflows: uncertainty propagation, which characterizes route-level score distributions under the retained Monte Carlo design, and sensitivity screening, which identifies which sampled foreground levers are most strongly associated with score variation. Reported results were presented as ranges rather than point predictions. In the case study, a subset of key results was carried through each step so that the implications of specific choices could be traced from the baseline LCA to the uncertainty analysis.

Operationally, the framework unfolded in five steps (Figure 3-1), which implemented that logic end-to-end:

> ***Step 1 — Establish Accounting Scope and Baseline Conventions:*** Goal and scope (per ISO 14040/44), cut-off criteria, and allocation via system expansion were defined. Temporal alignment was specified such that all consequences were booked in the deconstruction year, consistent with consequential practice (B. Weidema, 2003; Ekvall & Weidema, 2004). Stage D substitution was specified to satisfy like-for-like functional and quality equivalence, with equivalence modelled explicitly prior to crediting. In the case

study, the admitted EoL route register was also fixed here as a scope input to later market and substitution steps. A contemporary-baseline database snapshot was instantiated and frozen for all downstream comparisons.

Step 2 — Construct Scenario-Consistent Backgrounds: IMAGE pathways SSP1–VLLO, SSP2–M, and SSP5–H were processed with premise into ecoinvent 3.10.1 consequential as single-year assessment-year snapshots, focusing on scenario-consistent transformation of electricity supply and fuel/energy-carrier markets (Stehfest et al., 2014; Sacchi et al., 2022). These pathways operationalized scenario uncertainty for the prospective analysis. The contemporary baseline background remained the frozen "as-is" snapshot instantiated in Step 1.

Step 3 — Define Market Boundaries and Compute Marginal Supplier Mixes: Functional and market boundaries were operationalized for consequential linking within each background, with cross-border flows handled through market routing and trade exposure as required by the background data. A short multi-year trend window identified suppliers with sustained capacity change only where a foreground marginal-supply representation was required for credited product markets and where database resolution was insufficient for the Canadian decision context. Where premise/ecoinvent already encoded scenario-year commodity mixes (notably electricity and fuels, and hydrogen where represented), those mixes were inherited rather than recomputed. The resulting marginal supplier mixes formed the basis for Stage D substitution in the subsequent steps.

Step 4 — Model Foreground EoL Pathway, Stage D Substitution, and Deterministic Results: Concise unit-process models for the EoL pathways admitted in the application (as fixed in the case-study scope) were developed with transparent TRL references and explicit parameters (efficiencies, yields, and losses). These foreground models were linked to the frozen contemporary background and to scenario-consistent prospective backgrounds, using the marginal supplier mixes from Step 3. Like-for-like substitution with quality matching was enforced prior to crediting. Where equivalence was not met, routes were downscaled or excluded, and any required conditioning burdens (for example, purification, alloying, drying) were included before crediting (Arvidsson et al., 2018, 2023). This step produced deterministic LCA results for the contemporary baseline and a defined prospective case set across the admitted EoL routes. Results were decomposed to show the roles of background change, marginal supplier mix, and foreground modelling, and were presented with consistent visuals (stacked contribution charts and selective waterfalls) to aid interpretation.

Step 5 — Propagate Uncertainty and Test Sensitivity: Building on the Step-4 deterministic results, uncertainty was propagated selectively according to what was implemented and supportable in each time frame. For the contemporary baseline, only

background exchange uncertainty was sampled, while foreground route parameters were held at their appendix-defined central values. For the prospective scenarios, foreground parameter uncertainty was propagated within each frozen SSP background, and scenario uncertainty was represented through comparison across SSP1-VLLO, SSP2-M, and SSP5-H. Prospective background-exchange uncertainty and prospective joint background-plus-foreground propagation were not executed in the final study. Convergence checks and targeted sensitivities (e.g., trend-window length, hotspot overrides, conditioning choices) were performed. Results were reported as p5–p95 intervals within each scenario (Lloyd & Ries, 2007; Arvidsson et al., 2018; Cox et al., 2020; Mendoza Beltran et al., 2020). Sensitivity analysis was then performed as a separate global foreground screening exercise within fixed pathway backgrounds.

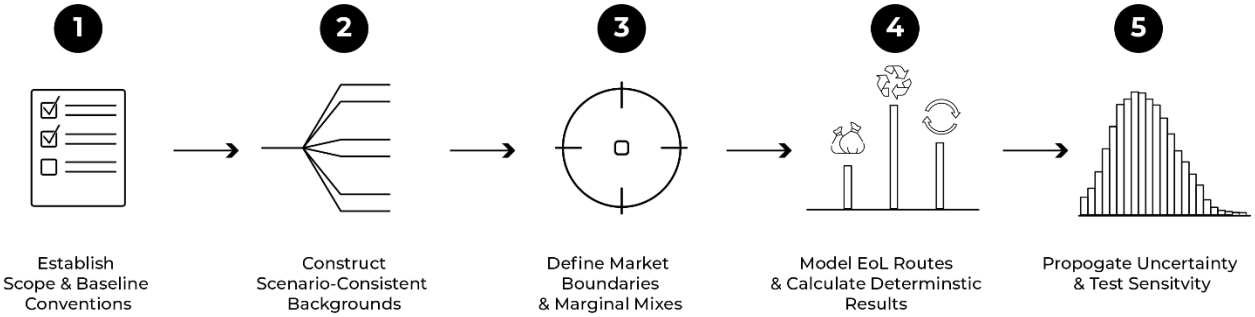


Figure 3-1: *Proposed workflow for prospective consequential LCA (pCLCA) of end-of-life pathways for the built environment: baseline accounting conventions, IAM-derived background scenarios, marginal-mix computation, deterministic Stage D substitution results, and uncertainty/sensitivity analysis*

3.2 Step 1: Establish Accounting Scope and Baseline Conventions

3.2.1 LCA Goal Definition & Objectives

The goal was to evaluate the environmental consequences of contemporary and prospective EoL strategies for curtain wall systems in a pCLCA framework. The LCA quantified current impacts and projected changes under evolving technological, market, and policy conditions to compare trade-offs across scenarios.

Two baselines were analysed, distinguished by booking year: *Contemporary*, representing present-day material flows, energy mixes, and prevailing EoL practices; and *Prospective*, reflecting projected background conditions with scaled adoption of emerging EoL technologies and market dynamics. The workflow linked evolving background systems with EoL-focused foreground modelling (C3–C4) and beyond-boundary substitution (Stage D) defined in later steps.

3.2.2 LCA Scope & System Boundaries

The study followed ISO 14040/44 and BS EN 15978-1 (ISO, 2006a, 2006b; BSI, 2026) and quantified EoL modules C3–C4 plus the beyond-boundary credit (Stage D). A1–A5 (upfront), B1–B7 (use), and C1–C2 (deconstruction logistics) were not quantified. This scoping reflected that (i) for an existing asset, upfront emissions were sunk and only weakly sensitive to EoL decisions; (ii) use-stage modules B1–B5 (maintenance, repair, replacement, refurbishment) were excluded because their timing and intensity were highly asset-specific, difficult to predict, often a smaller portion of overall emissions, and major replacements were, in practice, equivalent to component end-of-life choices (C3/C4 + D); (iii) operational energy (B6) was treated as an exogenous decarbonization pathway rather than a decision lever in this study; and (iv) B7 (operational water) was typically minor, poorly reported, and outside the main drivers examined here. Emphasizing C3–C4 + Stage D focused the analysis on the under-studied intersection of EoL pathway uncertainty and prospective substitution effects central to the research questions (Figure 3-2).

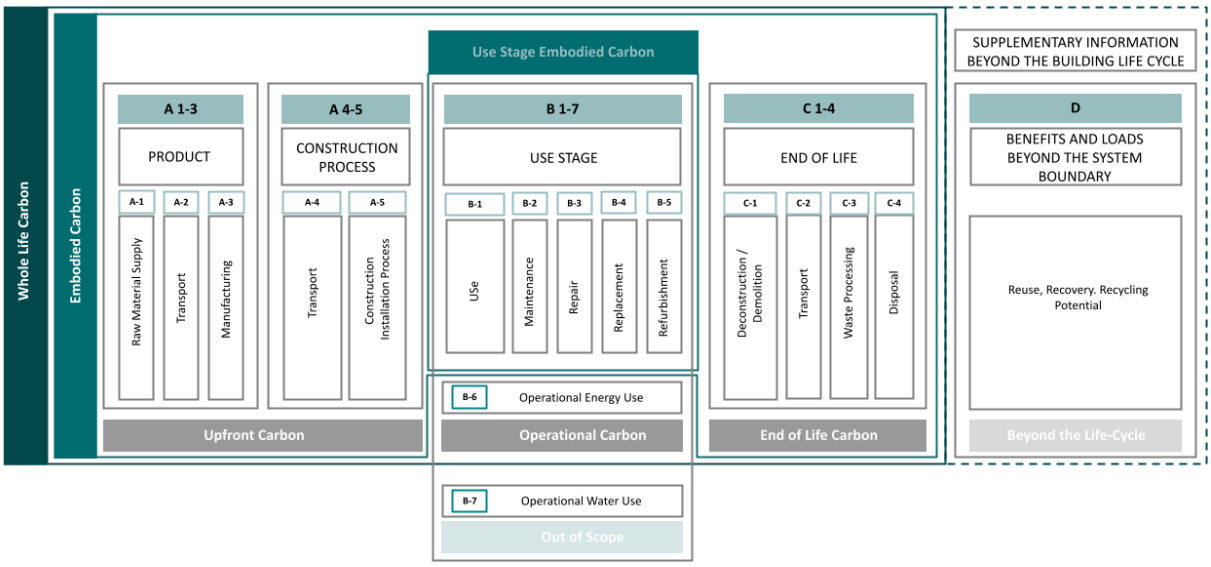


Figure 3-2*: Whole-life carbon assessment scope and module structure. Life cycle modules are shown following the building life cycle module convention (A1–A3 product stage, A4–A5 construction, B use stage, C end-of-life, and Module D beyond the system boundary), with included and excluded modules delineated by the system boundary.*

3.2.2.1 Accounting Conventions

- ***Geographical Scope*****:** The decision context was Canada, with regionalized data for electricity mixes, material sourcing, and waste-management practices. Under the consequential boundary, markets were followed where they operated. When Canadian demand changes triggered supply adjustments abroad, international trade flows (imports/exports) were included and displaced products credited in their actual geography, using region-tagged processes and market boundaries.
- ***Temporal Representativeness*****:** Two booking years were evaluated, aligned with the *Contemporary* baseline and *Prospective* baselines. All consequences were booked in the deconstruction year so that EoL burdens and any Stage D credits reflected background

conditions when they occurred. Prospective scenario definitions were introduced in Step 2, and Stage D implementation details were detailed in Step 4.
- ***Cut-off Criteria*:** No a priori exclusions. After results were generated, contributors ≤ ~1% of total impact could be pruned from reporting. Any pruned items were listed in an appendix.
- ***Allocation*:** Multifunctionality was handled via system expansion and partitioning was avoided unless required by data constraints and documented when used.
- ***Data Sources*:** Ecoinvent v3.10.1 consequential, regionalized for Canada, was used and supplemented by expert judgment where coverage was limited. A contemporary baseline snapshot was frozen for quality assurance and downstream comparisons.

3.2.2.2 Admission Policy (applied to emerging routes only)

To ensure that any emerging EoL pathway carried forward to foreground modelling and Stage D substitution was both auditable and consequentially interpretable, an emerging route was admitted only if it satisfied both conditions below:

1. ***Traceable LCI Parameterization***: The route could be represented as an LCI-compatible unit-process chain using directly reported data, documented mass/energy balances, and/or a justified proxy transfer that preserves physical coherence. Black-box placeholders without auditable parameterization are excluded.

2. ***Plausible Substitution Intent***: The route's output could be stated as a plausible displaced-product intent (form and indicative specification class) within the consequential framing, sufficient to define a candidate displaced market in Step 3. Formal equivalence qualification (like-for-like checks, required conditioning, and any downscaling or exclusion) is deferred to Step 4 when Stage D is implemented.

Application-specific route selection is reported in Chapter 4. The full longlist, exclusions, and screening notes from the targeted review were archived in Appendix B (as appropriate).

3.3 Step 2: Construct Scenario-Consistent Backgrounds

This step defined how prospective background databases were built and made consistent with an internally coherent future system state. Scenario pathways from an Integrated Assessment Model (IAM) were translated into life cycle inventory (LCI) form and applied to an ecoinvent consequential database for a future reporting year. The result was a "frozen" background database representing one internally consistent future state. That frozen background was then used, unchanged, for all marginal supplier analyses and consequential modelling. The transformation was implemented as a single-year operating-state snapshot for the assessment year. Multi-year IAM trajectories served as upstream evidence for interpolation but were not treated as multi-year inventories in the LCA.

Consistent with recent pLCA guidance, Step 2 was limited to constructing and freezing the scenario-consistent background for the assessment year. Because the prospective backgrounds were generated in an ecoinvent consequential context using premise, Step 2 also fixed the

pathway-specific electricity and fuel/energy-carrier markets consumed by the rest of the background system, including the marginal-supply representation embedded in those energy markets under the selected premise configuration. Marginal suppliers for decision-linked credited product markets (e.g., Stage D displacement markets and any non-energy markets requiring additional regional resolution) were then identified in Step 3, where market boundaries and trend windows were operationalized explicitly as a foreground construct. Tool-derived marginal mixes were adopted in Step 3 only where their market definitions and parameters aligned with the credited product market; otherwise, foreground rules were applied (Sacchi et al., 2022; Maes et al., 2023, 2024).

This aligned with best practice in prospective LCA, which stressed (i) temporal alignment to the decision year, (ii) internal scenario consistency across sectors, and (iii) reproducibility of the background system.

3.3.1 Scenario Selection and Bounding Uncertainty

To span structural (scenario) uncertainty in the prospective backgrounds, three contrasting pathways were used—SSP1–VLLO, SSP2–M, and SSP5–H (Maes et al., 2023; Šimaitis et al., 2025)—following pLCA guidance to cover key drivers across energy, industry, and materials (Lueddeckens et al., 2020; Thonemann et al., 2020):

- ***SSP1-VLLO*:** Represented a sustainability-focused, optimistic pathway with very-low radiative forcing and low/no overshoot. It assumed rapid decarbonization, strong efficiency gains, accelerated electrification, and coordinated policy—consistent with keeping end-century warming in the low range.

- ***SSP2-M***: Represented a "middle-of-the-road" pathway with medium forcing. Socioeconomic trends were mixed, climate policy was moderate, and technology improved steadily but not dramatically.

- ***SSP5-H***: Represented a pessimistic, fossil-fuelled development pathway with high forcing. It featured energy-intensive growth, continued reliance on fossil fuels, and weaker or delayed climate policy, leading to higher end-century climate warming outcomes.

3.3.2 Prospective Pathway Integration

For each selected pathway IMAGE provided the projected energy system and associated technology configuration for the prospective assessment year, including electricity generation technologies, operating efficiencies, direct emissions, carbon capture assumptions, and the production routes and carbon intensities of major energy carriers such as natural gas, refined petroleum products, biofuels, hydrogen, and synthetic fuels (Stehfest et al., 2014; Van Vuuren et al., 2014; Harmsen et al., 2021). The software premise ingested these IMAGE projections and generated corresponding market activities in ecoinvent for that pathway and year (Sacchi et al., 2022).

In this study, only the electricity supply system and the fuel/energy carrier supply system were structurally transformed in the prospective backgrounds. For each of the three pathways, premise generated pathway-specific electricity and fuel markets from the IMAGE scenario drivers as a single-year assessment-year snapshot using the premise "foresight" (PERF) configuration, and these markets were inserted into ecoinvent to replace the default suppliers for the given pathway. In the consequential database context, this choice fixes the pathway-year energy market composition that upstream processes draw from, including the marginal-supply representation encoded in those electricity and fuel markets. All applicable activities in the prospective background were then linked to these pathway-specific electricity and fuel markets rather than to contemporary energy suppliers. As a result, every activity that consumed electricity and/or fuel in each prospective background operated under the energy system defined by that IMAGE pathway for the selected assessment year.

Cement production, steel production route shares, freight transport, industrial heat supply, and engineered carbon removal were additional sectors that could reflect a future state as defined by IMAGE. However, these sectors remained in their original structural configuration. This was because these sectors were not accompanied, at the required regional and consequential resolution, by a pathway-specific long-run marginal supplier mix that could be applied consistently in this study's marginal displacement logic (Sacchi et al., 2022; Maes et al., 2023). These sectors were therefore evaluated under the updated prospective electricity and fuel supply for each pathway but were not otherwise redefined, and any unmodified sectors were documented as gaps rather than partially rebuilt. This reflected a current scope limitation of IAM-to-LCI pipelines. Accordingly, Step 2 was explicitly an "energy-first" scenario translation, and it ensures scenario-consistent energy and carrier markets while avoiding partial, non-reproducible rebuilding of other sectors.

3.3.3 Data Processing & Representation of Scenario Drivers via Premise

Post-processing was intentionally limited to deterministic QA and reproducibility controls rather than substantive manual edits to the background inventories. For each scenario-specific background, the following checks were executed and recorded:

- ***Biosphere Consistency:*** All biosphere exchanges pointed to the project's ecoinvent-aligned biosphere database.
- ***Link Integrity:*** There were no broken technosphere links within the background database and there were no orphan references to non-existent databases.
- ***Basic Sanity Checks:*** Activity counts within expected ranges and successful matrix processing.
- ***Reproducibility Controls:*** Database naming encoded pathway/year/foresight mode, environment versions (ecoinvent system model, premise version) were recorded, and a deterministic checksum over activity identity fields was stored in a manifest.

After these checks, each background database was designated as frozen (recorded via manifest and metadata markers) and used unchanged for all subsequent modelling steps.

The prospective backgrounds retain the IAM's native regionalization and ecoinvent's geography structure. Canada-specific granularity is not introduced at the background level. Canada-specific consequential interpretation (e.g., marginal suppliers for credited markets, Stage D displacement definitions, and any jurisdiction-specific foreground substitutions) was therefore implemented as a foreground construct in later steps.

Further detail on premise's internal mapping logic and interpolation procedures is provided in Sacchi et al. (2022).

Where a sector's tool-derived marginalization settings (time window, foresight/lead times, lifetime/capital replacement, growth metric) and its market boundary (product grade/spec and geography) coincided with the credited product market defined in Step 4, those tool outputs could be adopted directly. Otherwise, marginal mixes were identified transparently in the foreground using functional-equivalence and grade/spec rules, supported by Step-2 CSV evidence and trade-routed exporter bins (per the heuristic in §3.5).

3.3.4 Limitations of Scenario Translation

Limitations remain due to regional aggregation, sector coverage (electricity/fuels focus on consequential premise workflows), and the need for proxies where IAM-to-LCI mappings are incomplete. These were treated as structural/model uncertainties and were discussed in Chapter 2. Decision-relevant uncertainties were explored via targeted sensitivities in the case study rather than by introducing non-reproducible manual changes to the background system. Because premise outputs are typically regional (not subnational), within-Canada heterogeneity is not represented at the background level. Any subnational specificity required for foreground operations was implemented explicitly in later steps as a documented modelling choice.

3.4 Step 3: Define Market Boundaries and Compute Marginal Supplier Mixes

Ecoinvent (consequential) provided a consistent background for this study. However, at the inspected resolution, several hotspot markets still required additional Canadian foreground detail as supplier shares were sometimes available only at broader regional levels, Canada-specific splits were not always reported, and data were not consistently available as continuous annual time series. For these reasons, marginal suppliers were identified in the foreground.

Marginal suppliers were defined as the producers or technologies expected to adjust output in response to incremental demand, consistent with CLCA practice (Weidema et al., 1999; B. Weidema, 2010). The approach applied to both contemporary and prospective scenarios while remaining aligned with the Canadian decision context. In all cases, Step 2 provided scenario-conditioned inventory contexts, while Step 3 defined operational market boundaries and

computed marginal mixes as a foreground construct. Computational details and worked examples are provided in Appendix D. This section set out the governing policies.

3.4.1 Step 3-1: Identifying the Scale and Time Horizon of the Change

Changes were treated as incremental relative to market size. In the contemporary baseline they occurred within existing market boundaries. In prospective scenarios they acted within structurally evolving markets. Supplier identification remained a foreground construct and the Step 2 scenario-conditioned inventories provided techno-environmental context only.

The response horizon was four years: a trailing window up to the booking year for the contemporary baseline (e.g., 2017–2020 for a 2020 booking), and a centred window around the booking year for prospectives (e.g., 2029–2032 for a 2030 booking) (B. Weidema et al., 2013; Maes et al., 2023). Where data were incomplete, a contiguous ≥2-year window was used and the justification logged. No look-ahead bias was introduced in the contemporary case. Where Canadian data availability or sector reporting made the four-year horizon non-operational for a hotspot market, a longer fixed window could be adopted as a documented exception, provided it was applied consistently within that market and not mixed across scenarios.

Where credible forward-looking evidence exists for long-run response (e.g., published electricity supply projections), forecasted trends may be used in preference to purely historical trend inference, consistent with consequential guidance that forecasts were preferable when available. In such cases, the forecast horizon was treated as the response horizon for marginal identification and is documented as a market-specific exception.

Where marginal mixes were inherited directly from the premise-generated IAM scenario database for a given year (e.g., electricity and fuels in IMAGE 2050), the effective identification interval followed the consequential settings embedded in premise (e.g., range time and foresight configuration) rather than an externally imposed centred window.

Where prospective evidence was provided primarily as long-horizon growth factors rather than annual series, those factors were converted to implied annualized growth rates and applied consistently within the selected response horizon. These trends were treated as scenario inputs rather than observed time series.

3.4.2 Step 3-2: Identifying the Limits of a Market

Markets were defined by functional equivalence, grade/specification, geography, and trade exposure, following established consequential guidance as outlined by Weidema (2003). Functional and specification limits were applied first. Products failing equivalence screens (function, grade, or specification) were excluded from the candidate market.

Geography was defined operationally in a manner consistent with the Canadian decision context. For the contemporary baseline, procurement practice and observed supply channels governed market limits. For prospective scenarios, the same market-definition rules were applied while technology availability and upstream conditions were inherited from Step 2.

For trade-exposed hotspot markets, geography was operationalized as Canada plus importer bins representing the contemporary supply regions feeding the Canadian market. Where Canadian trade reconciliation required accounting for net exports to balance supply and demand, exports were handled at the reconciliation stage of the Canadian market representation rather than treated as explicit Canadian "exporter bins" within the marginal mix. Outside predefined hotspot markets, ecoinvent consequential market definitions were accepted to preserve transparency and avoid ad hoc boundary changes.

3.4.3 Step 3-3: Identifying Trends in the Volume of a Market

For each hotspot market, defined in relation to the admitted EoL route set, a primary evidence source was selected to establish supplier composition and trend information. In the contemporary baseline, Canadian production statistics with reconciled trade balances served as the primary source. In prospective scenarios, the source was chosen case-by-case to balance alignment with the scenario-conditioned inventories and fidelity to the Canadian decision context. Geography was Canada by default and was scaled to North America only where Canadian-level resolution was unavailable or non-representative for projections.

For prospective markets, the material-specific hierarchy was:

(i) IAM-derived series (e.g., IMAGE) where sectoral mapping to Canadian markets was robust (e.g., electricity, fuels);
(ii) non-IAM, Canada-anchored outlooks and harmonized trade or capacity projections where IAM mapping was weak or non-representative for the market at the required resolution;
(iii) supplier-set–constrained prospective modelling for traded commodities where future trade reconfiguration and new entrants could not be defensibly forecast at the required resolution. Under this approach, the contemporary supplier set was retained as the defensible "who supplies?" structure, while prospective change was introduced using region-level or technology-level trend information from reputable long-horizon outlooks and sector roadmaps (e.g., IEA and industry or consortium pathways). Where appropriate, these signals were used either to derive implied growth within the selected response horizon or to apply technology-evolution overlays to the emissions intensity of supply, without inventing facility-level forecasts or speculative reallocation of supplier shares;
(iv) carry-forward of the 2025 marginal composition where no suitable prospective series existed, flagged for sensitivity in § 4.7.
(v) and, lastly, Canada-anchored technology-evolution overlays may be applied in prospective scenarios when IAM mapping is weak at the required resolution, but credible Canadian evidence exists that the production mechanism itself will change

(e.g., primary aluminium anode technology). In these cases, supplier shares are treated as a "who supplies?" construct, while the technology overlay modifies the "how emissions-intensive is supply?" construct without inventing facility-level forecasts.

For prospective scenarios, IAM- or outlook-based projections used under tiers (i) and (ii) were treated as authoritative scenario inputs and were not adjusted through additional plausibility screening or post hoc alignment. Under tier (iii), by contrast, the retained supplier-set structure was not interpreted as a forecast of future marginal market shares; rather, it provided a constrained decision-context baseline onto which defensible prospective growth or technology signals could be applied. This distinction was especially important for traded commodities where credible plant-level or importer-specific forecasts were unavailable.

Tier (iv), defined as carry-forward of the most recent defensible marginal composition, was treated as an explicit data-limitation sensitivity rather than as a probabilistic uncertainty distribution. Where tier (iv) was invoked in prospective scenarios, the carry-forward mix was reported as the central default and paired with discrete alternative cases, such as conservative and aggressive shifts informed by the nearest available regional or technology signals, or an extreme-bound case that reassigned marginality to the dominant contemporary bin.

For the contemporary baseline, Step-4 marginal mixes were derived from observed Canadian production and reconciled trade statistics over the trailing window and were not subjected to market-mix uncertainty propagation in this thesis. Where uncertainty analysis was performed, it addressed foreground process parameters and other specified sources (see §4.7), not the contemporary marginal mix definition.

3.4.4 Step 3-4: Identifying Changes in Supply & Demand

For each hotspot market, the co-marginal supplier set was defined as those suppliers that exhibited sustained growth over the analysis window. When no supplier grew, the set consisted of those with the least decline. Shares within the marginal mix were apportioned deterministically and remained consistent with the market limits and the selected evidence source.

Shares were computed using the growth signal available at the resolution of the evidence source:

(i) where supplier volumes were available as window endpoints, shares were proportional to positive absolute output change over the window ($\Delta V > 0$), normalized within the co-marginal set;
(ii) when all candidates showed non-positive (declining) trends, shares were proportional to the baseline composition from the same source family selected in § 3.5.3 and restricted to the co-marginal set;

(iii) when growth data were incomplete for a supplier, the nearest contiguous ≥2-year window (or the tier (iii) implied growth model for that supplier bin) was used and flagged.

Where forecast-based evidence is used (e.g., electricity projections), the same co-marginal logic is applied to the forecasted change signal over the forecast horizon. Shares were normalized to unity and recorded. The procedure remained a foreground computation and did not alter the frozen Step-2 backgrounds. Final marginal mixes were versioned by scenario-year and logged in Appendix D.

3.5 Step 4: Model Foreground EoL Pathways and Stage D Substitution

This step implemented the admitted route register and case-study inclusion choices (reported in Chapter 4). It specified the foreground EoL models on top of the frozen backgrounds (Step 2) and the marginal mixes (Step 3). For each admitted route, the step defined the unit-process chain, the equivalence checkpoints, any required proxies, and the Stage D mapping. The objective was to represent conventional and emerging options at unit-process resolution and to apply Stage D credits only after like-for-like substitution. Inclusion choices from §3.4 were not revisited. Case-specific route selection, technology maturity characterization, and excluded candidates are reported in the application chapter and its appendices. The methods here specify only the modelling rules.

This step produced deterministic LCA results for the contemporary baseline and a single central prospective scenario set across all candidate EoL routes. Results were decomposed to show the roles of background change, marginal supplier mix, and foreground modelling, and were presented with consistent visuals to aid interpretation. Additionally, the deterministic outcomes of Step 4 defined the single "central" parameter set drawn from the evidence-based ranges established for Step 5 (i.e., a point estimate that is explicitly consistent with the later uncertainty model), while all interval reporting is reserved for Step 5.

3.5.1 Foreground EoL Route Modelling

Foreground models were constructed for all EoL routes (defaults and emerging) as concise unit processes with explicit inputs, utilities, yields, and losses. Deterministic links captured dependencies among these parameters (e.g., yield–energy and recovery–loss trade-offs). Inputs and utilities were sourced from the Step 2 backgrounds. A standardized "route wrapper" was used, but the internal normalization was defined as a route-entry reference flow that matches the first modelled transformation (e.g., 1 kg of recovered product). A deterministic conversion map (yields, metal fractions, and loss factors) carried that route-entry basis through the chain and back to the study's common accounting basis (and ultimately the functional unit (FU)), ensuring comparability without forcing non-physical unit conversions at the unit-process level. The wrapper was then scaled to the functional unit (FU) to assemble the overall LCA.

3.5.2 Proxy Foreground Processes

Proxies were used where inventories did not exist in ecoinvent for required sub-operations, and where emerging technologies required unit operations that are not represented at commercial scale in the database. Proxies were implemented by cloning the closest available activity and applying a minimum-change rule: only exchanges needed to represent the new unit operation were edited (reference product, key inputs/yields, and key utilities), while the remaining exchange structure was retained to preserve upstream realism.

To keep proxies auditable and reproducible, each proxy was required to:

(i) maintained closed mass and energy balances for the edited transformations,
(ii) exposed all edited elements as an explicit parameter list (units, central values, bounds, and provenance), and
(iii) linked deterministically to adjacent foreground steps to prevent burden "leakage" into background markets.

Because several sub-operations are novel and data-limited, uncertainty was assigned primarily to the parameters that define the edited behaviour of the proxy or emerging step (e.g., yield/loss fractions, recovery, and specific energy), while unedited exchanges retain the uncertainty structure of the underlying database. Proxies were then connected to each route through explicit demand quantities computed by the route's deterministic mass balances. Appendix B reports proxy definitions, parameters, and sources. Where a proxy's native unit differs (e.g., m^3 vs kg), a documented stoichiometric or physical conversion was applied to express the proxy per unit of required service before linking.

3.5.2.1 Utility Regionalization of Foreground Proxies

When a suitable base activity existed but not in the desired geography, a targeted geographic adjustment was applied. In this work, this was best described as utility regionalization where proxy electricity, fuel, and district heat/steam providers were relinked to geography-consistent markets (and to the study's marginal electricity construct where applicable), while the rest of the process structure was left unchanged.

This scope is justified because energy carriers were among the most geography-sensitive drivers in unit-process LCIs (Sacchi et al., 2022), and it aligns with the study's broader prospective setup, where scenario coupling primarily updates electricity and fuels. Remaining non-utility geographic differences were treated as residual background/model uncertainty rather than being selectively edited without strong evidence.

3.5.3 Evidence-Based Scaling

Evidence-based scaling was used to translate heterogeneous evidence on emerging EoL technologies into implementable foreground inventories that remain traceable, physically coherent, and comparable across routes. Because emerging technologies are rarely reported in a

form that maps directly to LCA unit processes, the objective was not to solely reproduce any one study, but to extract measurable, defensible operating parameters (for example, specific energy demand, material yield/loss factors, and material/chemical consumables), where available, and convert them into a consistent unit-operation structure compatible with the study's C3–C4 system boundary (Buyle, Galle, et al., 2019; Thonemann et al., 2020).

Sources were screened primarily for transferability to the modelled unit operation: whether the reported data corresponded to a comparable feed condition, operating regime, and measurement basis, and whether they reported the specific variables required to parameterize the inventory. Where direct transfer was weak or reporting incomplete, the missing elements were not absorbed as implicit assumptions. Instead, they were handled transparently by (i) using conservative proxies, (ii) bounding values using multiple independent references, and/or (iii) promoting the parameter to Step 5 as explicit uncertainty or a labelled sensitivity case.

To keep the baseline consistent across all emerging routes, Step 4 fixed a single set of deterministic central values for each technology under the contemporary background and the prospective backgrounds. Technology performance levers that are known to vary with configuration, scale, and operational control were reserved for Step 5 so they could be tested systematically rather than embedded as unexamined improvements (Thonemann et al., 2020; Tsoy et al., 2020). Importantly, the Step 4 values are deterministic reference configurations used to define the baseline route implementations for later comparison. They informed the selection and centring of the retained uncertainty levers, but they were not interpreted as quantities that must numerically equal the p50 of the uncertainty simulations. Step 5 was executed as a separate stochastic analysis on the implemented uncertainty layers, and because the reported net result is the residual of burden-side and credit-side terms under sampled parameters, the medians of those stochastic distributions did not need to reproduce the deterministic point estimates exactly.

For each included emerging pathway, a report-then-scale rule was followed. Processes were modelled at contemporary specification and then were scaled to prospective backgrounds using referenced improvements in yields, efficiencies, or loss reduction. Where evidence was thin, improvements were bounded conservatively and were flagged. Background changes propagated through the Step 2 databases. Scaling relied on the unit operations defined in §3.6.1–3.6.3, including proxies. No new, unverified operations were introduced during scaling.

3.5.4 Stage D Substitution

Stage D substitution converts equivalent recovered outputs from the modelled end-of-life routes into avoided production within the scenario-year marginal supplier mixes described in §3.5 and operationalized in §4.5. Consistent with EN 15978-style "beyond-the-system-boundary" accounting, these effects were implemented as an external credit rather than as an in-system allocation. In practice, Stage D benefits were treated as conditional consequences of end-of-life treatment: they were applied only when the modelled route yields an output that can plausibly

displace an incumbent market product in the receiving system, after accounting for all conditioning burdens and losses required to reach that product's specification.

To implement this consistently across routes, each end-of-life option was wrapped in a standardized route wrapper that exposes an explicit route-entry reference flow and a deterministic conversion chain to the creditable product(s), producing net recovered products for crediting. All conditioning burdens and losses were applied before any avoided-burden credit. Eligible quantities were then multiplied by the relevant scenario-year marginal mix and booked in the deconstruction year, ensuring avoided production reflects the active background at the time of displacement.

Product-to-market mappings, equivalence decisions, and any downscaling rules are reported in Chapter 4, with structured mapping tables archived in Appendix A. Where a route produced multiple eligible products, each product was assessed and credited independently once it met the equivalence requirements.

3.5.4.1 Equivalence & Conditioning

The equivalence and conditioning rules applied only to Stage D substitution. Conditioning was modelled in the foreground up to the point at which the recovered output met the displaced product's specification; only then was the output routed to the appropriate background market and credited as avoided production.

Equivalence was assessed via an explicit screen requiring alignment on:

a) State/form (e.g., billet vs flake; cullet class; pulp class),
b) Grade/specification (e.g., alloy/strength class; moisture/contaminant thresholds; fibre class), and
c) Market plausibility/geography consistent with the consequential boundary (i.e., the credited product is one that could credibly be supplied into—and displace within—the relevant market definition).

Only conditioning necessary to satisfy this equivalence screen was included. Typical conditioning operations included drying, sizing, finishing, purification, alloying, decontamination, and upgrading steps required for market acceptance. Any associated losses were applied prior to crediting, such that Stage D reflects net displaced quantity rather than gross recovered mass.

Implementation Rules:

- Credits were applied only after the equivalence screen is met
- All conditioning burdens and losses were included before crediting net displaced quantity.

- If equivalence was uncertain, credit was downscaled or excluded rather than assuming full displacement.

All foreground parameters and deterministic links were tabulated in Appendix B and instantiated in Chapter 4.

3.5.5 Deterministic LCA

Deterministic LCA results were presented for the contemporary baseline and all three prospective scenarios (SSP1–VLLO, SSP2–M, SSP5–H) using the scenario-consistent frozen backgrounds and marginal mixes defined in Steps 2 and 3. Impacts were quantified using ReCiPe 2016 Midpoint (H), with climate-change results reported using the method's GWP100 characterization. Results were shown with consistent visuals (i.e., stacked contribution charts) to aid interpretation. Deterministic results were used as (i) the reference point for subsequent uncertainty propagation within each prospective pathway.

3.5.5.1 Contribution Analysis & Hotspot Features

Deterministic screening of the baseline foreground models identified operational add-ons that could shift outcomes at scale. Examples include internal recirculation, off-gas capture, and enhanced conditioning. These elements were treated as features for optional targeted exploration in Step 5, while their baseline configurations remained in the inventory where structurally required or typical for reported practice.

An operational add-on was designated a feature when it:

- was discretionary or configurable rather than intrinsic to core stoichiometry or mass balance,
- showed high materiality in deterministic results, and
- had uncertain or variable prevalence/performance across plausible implementations.

Features remained in the inventories at defensible central settings where applicable, and no removals were made solely because a contribution was large. For uncertainty treatment, features were analysed through structured scenario runs and targeted sensitivity sweeps over justified ranges. Occasionally, scenarios used explicit on/off toggles or level settings with documented rationale rather than pooled into Monte Carlo distributions. Technology-specific feature definitions, bounds, and scenario matrices were documented in Chapter 4, with parameter sources and justification tabulated in Appendix B.

3.6 Step 5: Propagate Uncertainty and Test Sensitivity

Uncertainty was propagated across the implemented uncertainty layers only. For the contemporary baseline, background exchange uncertainty was propagated using the uncertainty information embedded in the background inventories, while foreground route parameters were held fixed at their deterministic appendix-defined central values. For prospective scenarios, foreground parameter uncertainty was propagated within each frozen scenario-year background,

and scenario uncertainty was represented through comparison across SSP pathways. Prospective background exchange uncertainty was not propagated, and prospective joint background-plus-foreground uncertainty attribution was not executed in the final implementation. The resulting uncertainty summaries should therefore be interpreted as stochastic descriptors of the implemented Monte Carlo analyses, not as numerical reproductions of the deterministic Step 4 point estimates.

3.6.1 Uncertainty Taxonomy

Uncertainty was treated as three layers:

- ***Background Exchange Uncertainty*** *(Contemporary):* Propagated via Monte Carlo on background exchanges (e.g., ecoinvent pedigree parameters) and reported as within-scenario variability. If exchange uncertainty parameters were not present/retained in the prospective database, background uncertainty was not sampled, and the background was treated as fixed by construction of the IAM-coupled inventory; this was recorded explicitly as a methodological limitation.

- ***Foreground Parameter Uncertainty*** *(Prospective)*: Propagated via Monte Carlo on emerging technology parameters and reported as within-scenario variability (interval bands). For the contemporary baseline, emerging-technology parameters were fixed at evidence-based central values (Step 4) to avoid introducing additional epistemic variability where values were explicitly specified.

- ***Scenario Uncertainty*** *(Prospective)*: Represented by repeating the foreground Monte Carlo within each prospective pathway (SSP1–VLLO, SSP2–M, SSP5–H), with the pathway background held fixed per Step 2.

3.6.2 Foreground Parameter Uncertainty

Foreground parameter uncertainty was applied only to parameters associated with low- or med-maturity. The intent was to represent uncertainty arising from limited empirical maturity, implementation variability, and scale-up dynamics, while avoiding double-counting by enforcing deterministic links wherever physical relationships are known (e.g., mass balance, reaction stoichiometry, bounded efficiencies, and yield–loss closure). Where dependent quantities were governed by explicit route equations, they were recomputed deterministically from the sampled primitives rather than sampled independently.

3.6.3 Evidence-Based Range Elicitation

Foreground distributions for low-maturity processes were constructed using a two-step rule. First, parameters were decomposed into (i) measured/base components and (ii) scale-up/implementation components. Second, distributions were assigned based on the evidence class available for each parameter and the physical constraints that must be respected.

To avoid incoherent sampling and double counting, parameters were classified by whether they represented physical relationships that should hold deterministically, or implementation-dependent performance that plausibly varied with maturity and deployment:

- ***Deterministic Constraints*** (not sampled): Conserved mass balances, closed split flows (summing to unity), and stoichiometric relationships used to compute by-product potentials and required reactants (where applicable). These were treated as deterministic mappings to prevent sampling dependent quantities independently.

- ***Bounded Performance Fractions*** (sampled): Yields, conversion fractions, recoveries, and purity fractions were treated as *bounded parameters* on $[0,1]$where they represent proportions or efficiencies.

- ***Positive Intensity Parameters*** (sampled): Quantities expressed as rates (e.g., kWh/kg, kg/kg, L/kg) were constrained to be strictly positive and were modelled with distributions that allow right-skew where appropriate (e.g., Lognormal/PERT).

Distribution families were then selected to reflect both the *evidence class* available for a parameter and its *mathematical form*. Where replicate measurements or tight empirical variation were reported, parameters were modelled with truncated Normal distributions (or Lognormal when the parameter is strictly positive and empirically right skewed). Where variability was driven primarily by implementation and maturity (for example, overhead persistence, recirculation/make-up fractions, or improvement factors), parameters were encoded using PERT/Triangular distributions when a defensible "most plausible" value could be identified, or Beta distributions when the parameter was naturally interpreted as a bounded proportion on $[0,1]$. Where evidence was sparse and only an order-of-magnitude plausible range could be defended, a bounded Uniform distribution was used and explicitly flagged as "sparse-evidence," with influence checked via a paired one-way sensitivity test. In all cases, physical feasibility constraints were enforced directly through bounds (e.g., fractions constrained to $[0,1]$; rate/intensity parameters constrained to > 0).

3.6.3.1 Edge Cases and Sparse-Information Rules

Parameters lacking credible empirical support were not randomized by default. Instead, they were fixed at the best-supported central value and probed through explicit one-way sensitivity tests using wide but defensible bounds (typically ±50% unless constrained otherwise). If a parameter remained stochastic despite sparse evidence, a bounded uniform distribution was applied and explicitly flagged.

All draws respected feasibility constraints (e.g., non-negative flows; efficiencies within $[0,1]$); physically infeasible draws were rejected and redrawn. Where a maturity-related structural behavior was better represented discretely than continuously (e.g., an operational failure mode that triggers reprocessing), it could be represented as a Bernoulli event with an evidence-based

probability and added-burden consequence, with the assumption recorded and interpreted explicitly rather than hidden inside continuous ranges.

3.6.4 Monte Carlo Conditioning

Monte Carlo budgets were determined using convergence diagnostics rather than fixed draw counts. Fixed random seeds were used for reproducibility within each run configuration. Within a given scenario–year, the background database and marginal mix selections were held constant while eligible parameters were sampled, so that within-run variability reflects the intended uncertainty layer rather than unintended mixing of scenario contexts.

Convergence was assessed using a quantile-vector RMSE stability diagnostic evaluated at periodic checkpoints. At each checkpoint, the algorithm computed a vector of empirical quantiles for the primary impact method (GWP100) across the three reported cases (C3–C4, Stage D), and net (joint). Convergence was defined as stability of this quantile vector between successive checkpoints. Early stopping was triggered when the normalized RMSE fell below a predefined tolerance for a predefined number of consecutive checkpoint evaluation, following a minimum-iteration warm-up period. Additional details are provided in Appendix F.

3.6.5 Hotspot Features via Sensitivity Analysis

Hotspot features for emerging technologies were identified using foreground-only sensitivity analysis conducted within each prospective pathway under a fixed scenario background. Sensitivity analysis served as a screening step to identify which uncertain foreground assumptions exert the greatest leverage on the net results and therefore warrant priority for (i) targeted data improvement, (ii) tighter evidence-based parameterization, and (iii) potential follow-on stress tests in future work. This analysis is best understood as a simultaneous global screening exercise rather than as a one-factor-at-a-time perturbation test. At this stage, hotspot identification was used to prioritize review rather than to introduce additional feature-enabled Monte Carlo configurations.

Candidate levers were restricted to low- and medium-maturity foreground degrees of freedom (e.g., yields, conversion fractions, recovery/purity performance, energy intensities, and eligibility or quality-gate parameters) and to any purpose-built proxy edits admitted under the inclusion rule in §3.4.2 (Step 3). Parameters linked by deterministic physical relationships (e.g., mass-balance closure, stoichiometric constraints, split fractions summing to unity) were not perturbed independently. Sensitivity was evaluated on the governing independent parameters, with dependent quantities recomputed deterministically within each trial to avoid incoherent sampling and double counting.

Parameter influence was quantified using global, rank-based sensitivity metrics (e.g., Spearman rank correlation between sampled parameter values and the resulting net score) and summarized using driver scatter plots for the highest-ranked levers. Where sensitivity indicated that an

assumption behaved more like a discrete configuration choice than a continuous parameter (e.g., eligibility thresholds, alternative substitution ratios, mode-switching logic), it was recorded as a candidate structural hotspot for targeted follow-on scenario testing, rather than being pooled into baseline uncertainty propagation in the present study.

3.6.6 Scenario Uncertainty and Reporting

Scenario uncertainty was represented by repeating prospective foreground-uncertainty propagation across SSP pathways with pathway backgrounds held fixed, and reporting cross-pathway spread alongside within-pathway uncertainty where completed. Pathway coverage and convergence validation status were reported in the Results section. Sensitivity experiments and optional features were reported as separate deltas relative to the baseline rather than pooled into Monte Carlo distributions. Axes, units, booking year, and sign conventions were kept consistent; Stage D credits were plotted as negative.

3.7 Summary

This chapter formalized a five-step prospective consequential LCA (pCLCA) framework that couples IAM-derived, scenario-consistent background databases (frozen per pathway) with a minimal-override, Canada-aware marginal supplier protocol and foreground end-of-life (EoL) models that enforce like-for-like Stage D substitution. Uncertainty was treated in two complementary layers: structural uncertainty across IMAGE pathways and parameter uncertainty in background exchanges (where available) and in low-/medium-maturity foreground processes. Deterministic links (e.g., mass balance, stoichiometry, closure of split flows) were used to preserve physical coherence and avoid double counting when propagating uncertainty.

The methodological contributions centred EoL as the primary lens for pCLCA, an area that remains underdeveloped relative to upstream prospective modelling, and provided an end-to-end workflow that connects scenario-consistent backgrounds to unit-process EoL pathway representations and conditional Stage D crediting. The framework operationalized (i) explicit conditioning burdens and losses required to reach like-for-like substitution, (ii) a transparent proxy construction protocol when inventories or geographies were missing, and (iii) evidence-based scaling of emerging EoL pathways such that maturation effects were represented explicitly rather than embedded as implicit assumptions. Canada/North America–specific mapping and targeted regionalization procedures were applied to improve decision-context fidelity while maintaining traceability and reproducibility.

Chapter 4 applies this framework to a Canadian curtain wall case study. The case study instantiates the functional unit, booking years, and scenario set. It specifies the focal markets and admitted EoL route set and implements the frozen backgrounds, marginal supplier mixes, and foreground EoL pathway models to quantify C3–C4 burdens and Stage D effects in contemporary and prospective contexts. Uncertainty and sensitivity analyses are then reported

consistently with the framework, and key drivers are interpreted to support decision-relevant conclusions.

Chapter 4

4 Case Study: Application to a Curtain Wall

4.1 Overview & Reader's Map

This chapter implements the proposed framework via a case study of a curtain wall. To preserve a clear throughline, the aluminium subsystem within the 1 m² of installed curtain wall was followed consistently across Steps 1–5, with other components summarized in the appendices for completeness. End-of-life modules C3–C4 and the beyond-the-boundary Stage D were quantified with all consequences booked in the deconstruction year under a like-for-like substitution rule. The admitted EoL pathway set for the case study was intentionally compact and is documented as an application choice (rather than as a standalone methodological step) and encompasses disposal, reuse, conventional recycling, and a bounded set of emerging recycling routes grouped by maturity.

Two booking years were used: 2025 (contemporary) and 2050 (prospective). For 2050, scenario uncertainty was represented with three IMAGE-derived backgrounds constructed with premise (SSP1-VLLO, SSP2-M, SSP5-H). These prospective backgrounds were constructed, QA-checked, and frozen as single-year 2050 snapshots (PERF configuration), and then treated as the source of truth for background processes without subsequent manual restructuring. Foreground unit-process models for conventional and emerging end-of-life (EoL) options were layered on top of the frozen backgrounds and parameterized with explicit priors. Background exchange uncertainty was propagated only for the contemporary case, whereas the prospective pathway databases were held fixed within each SSP and uncertainty propagation focused on foreground parameters and scenario spread.

The evaluation space was organized as a compact scenario matrix that crossed the set of frozen backgrounds with a set of five EoL pathways. The matrix design remained general by construction: backgrounds captured structural differences in supply-side conditions, while rows spanned disposal, reuse, conventional, and two bands of emerging recycling at increasing technology readiness. This chapter reports per-scenario results with uncertainty bands and then compares scenarios to explain drivers.

Reader's Map

- ***Step 1 (Accounting Scope & Conventions):*** The functional unit (1 m²), booking years (2025, 2050), quantified modules (C3–C4+ D), primary indicator (GWP100), and a concise bill of materials (BoM) were fixed, and a contemporary consequential background was frozen for reference.
- ***Step 2 (Scenario-Consistent Backgrounds):*** Three prospective backgrounds were constructed and frozen from IMAGE via premise, with versions and hashes logged.

- ***Step 3 (Markets & Marginals):*** Market limits were defined, and marginal supplier mixes were computed over windows consistent with each frozen background only where foreground marginal construction was required for the Canadian decision context, otherwise mixes were inherited from the frozen backgrounds.
- ***Step 4 (Foreground EoL & Substitution):*** The admitted route set was implemented as foreground unit-process chains, including specification-critical conditioning, and Stage D substitution was enforced under the like-for-like rule. Deterministic results were reported for the contemporary baseline and the defined prospective case set.
- ***Step 5 (Uncertainty & Sensitivity):*** Contemporary background uncertainty and prospective foreground uncertainty were executed and interpreted alongside prospective scenario spread; all results were reported as functional-unit-level ranges within the implemented uncertainty structure.

This separation made attribution transparent: background change (Step 2), marginal-mix choices (Step 3), and foreground modelling (Step 4) were isolated and then recombined under the accounting rules established in Step 1. The aluminium focal elements and its data artefacts were introduced where they first became decision relevant. All non-focal materials and processes used the ecoinvent/premise background representations and default Canadian EoL processes without additional hotspot overrides unless explicitly noted in the appendices.

4.2 Step 1: Accounting Scope and Baseline Conventions

4.2.1 Functional Unit

The functional unit (FU) was 1 m² of installed, stick-built mass-timber curtain wall, deconstructed and processed in booking year $t \in \{2025, 2050\}$, while meeting the original panel performance specifications. All EoL activities necessary to handle that 1 m², including specification-critical processing/conditioning and the generation of any substitutable outputs, were included. Results were reported as GWP100 [kg CO_2e m²] as the primary indicator. Other midpoint categories were computed and are documented in Appendix E. A labelled schematic of the FU is depicted in Figure 4-1.

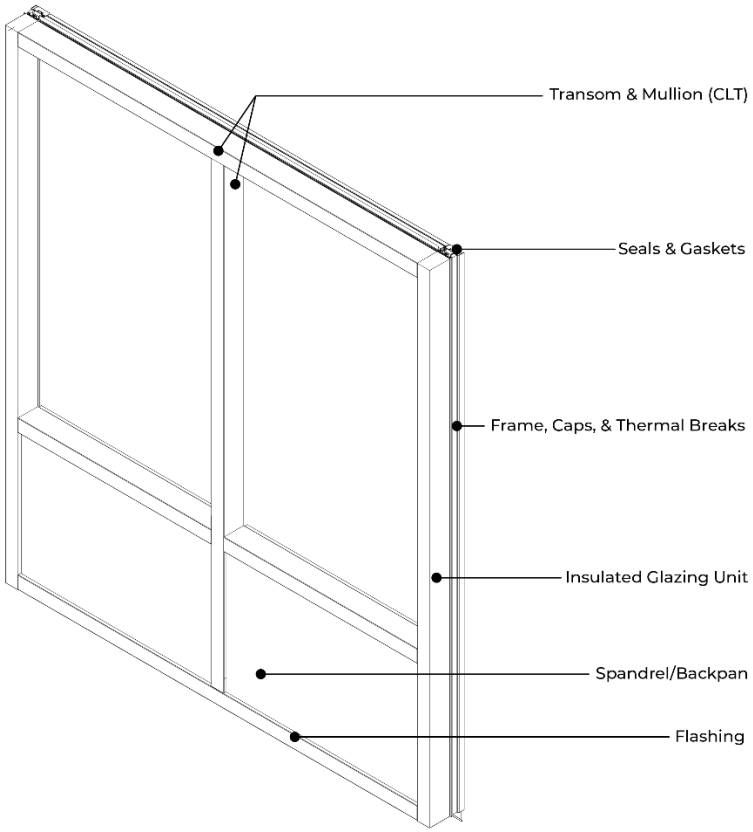


Figure 4-1*: High level schematic curtain wall panel with visible major assemblies labelled*

4.2.2 Scope & System Boundaries

The system boundary covered C3–C4 and Stage D. Upfront modules A1–A5 and use-stage modules B1–B7 were assumed to be held constant and thus not reported as they did not vary with EoL choices or influence comparative rankings. C1–C2 (site setup and dismantling/collection) were not quantified; standard collection practices and conversion rates were fixed across scenarios. Processing, conditioning, and disposal (C3–C4) were modelled as foreground unit processes when specification-critical (for example, alloy trimming and billet casting). Stage D crediting followed the like-for-like rule and was booked in the deconstruction year so that avoided burdens reflected the active background.

4.2.2.1 Geographic

The decision context was Canada. Consequential routing followed operational markets. When a recovered output displaced a product predominantly supplied by imports, the credit was taken in the exporter's geography using region-tagged processes rather than reassigned to Canada. Provincial electricity specificity mattered only insofar as it affected C3–C4 burdens and was encoded at the unit-process level. Displaced production inherited the geography of the relevant market defined in Step 3.

4.2.2.2 Temporal

Two booking years were used: *2025* (contemporary baseline) and *2050* (prospective). The 25-year interval approximated a single replacement cycle for curtainwall systems (typical service life ≈ 30 years). All burdens and credits were booked in the deconstruction year to avoid mixing contexts. For 2050, three separate background databases were derived from IMAGE via premise. Cross-year averaging was not used in inventories: when the booking year was 2050, inventories came exclusively from the 2050 background. The only use of short multi-year windows occurred in Step 3 for marginal-supplier detection and did not alter the inventory year for impact assessment.

4.2.3 Bill of Materials

A concise BoM for 1 m² anchored mass conservation across C3–C4 and Stage D and is summarized in Table 4-1 (full details are provided in Appendix B). Where specification uncertainty existed, acceptable alternatives were bounded using expert judgement and prior project experience; central values were recorded in Table 4-1. Quality variability of recovered materials at EoL was acknowledged but treated as out of scope. The analysis focused on EoL treatment pathways under the assumption of like-for-like equivalence, or a clearly stated functional proxy, for substitution credits. The BoM covered the timber frame, glazing unit, metals (anchors/fasteners aluminium profiles), and insulation as well as materials detailed in Table 4-1.

Table 4-1*: Curtain wall bill of materials aligned to major assemblies shown in Figure 10. Components are organized by assembly/zone (vision glazing, seals/gaskets, frame and thermal break, spandrel/backpan zone, flashings/trim, and subframing/backing) with material designation and quantity reported per square metre of installed curtain wall (kg/m²).*

Assembly	Component	Material name	Quantity (kg/m²)
Vision Glazing (Insulated Glazing Unit)	Glass	Double Glazed Flat/Float Glass	15.855
	Gas Infill	Argon	0.017
	Spacer	Aluminium	0.227
	Desiccant	Silica Gel	0.112
	Sealant	Silicone	0.085
	Setting Block	Silicone	0.281
	Glass Carrier	Fibre Glass	~0.001

Seals & Gaskets	Exterior Gasket	Ethylene Propylene Diene Monomer (EPDM)	0.217
	Mullion/Transom Gasket	EPDM	0.767
Frame, Caps, and Thermal Break	Transom & Mullion	Cross-Laminated Timber (CLT)	8.504
	Pressure Plate	Aluminium	1.179
	Front Face	Aluminium	2.175
	Thermal Base Profile	Polyamide	0.461
	Insulating Block	Polyamide	1.251
Spandrel / Backpan (opaque)	Backpan/Interior Finish	Gypsum	2.057
	Membrane	Polyethylene	0.056
	Backpan Insulation	Fibre Glass	0.168
	Insulation	Extruded Polystyrene (XPS) Insulation	0.132
	High Density XPS Block	XPS Insulation	0.008
	Coating	Metallic Oxide	N/A
Flashings & Trims	Horizontal Flashing	Aluminium	0.0530
	Vertical Flashing	Aluminium	0.0310
Subframing & Blocking	Framing	Softwood	0.359
	Blocking	Softwood	0.760
	Plywood	Softwood	1.452
Covers / Trim	Snap Cover	CLT	1.454

4.2.4 Admitted Route Register and Emerging-Technology Selection

This case study follows the aluminium subsystem of the curtainwall as the focal element because primary aluminium production and aluminium semi-fabrication remain a strategically important

Canadian manufacturing and export sector. Additionally, aluminium EoL pathways create large potential leverage in Stage D outcomes under a like-for-like substitution rule. Accordingly, the admitted EoL route register was defined explicitly as a Step-1 case-study scope choice so that Step 3 could unambiguously define displaced markets and marginal supplier mixes for the routes actually evaluated.

Rather than performing a full BoM-wide hotspot contribution screen, a targeted aluminium-only review was used to select a bounded set of emerging aluminium EoL pathways with high methodological relevance to prospective consequential substitution. The objective was to admit only those routes that could be represented as auditable unit-process LCIs and mapped to a clearly stated displaced-product intent, while deferring like-for-like qualification (equivalence thresholds, required conditioning, and any exclusion) to Step 4 when Stage D substitution is implemented.

4.2.4.1 Admitted Aluminium End-of-Life Routes

Defaults (boundary completeness / reference set):

- ***Landfill*** (no Stage D credit)
- ***Direct reuse*** of façade components (displaces like-for-like aluminium façade components, subject to later equivalence checks).
- ***Conventional recycling*** via remelting (closed-loop intent; displaces extrusion billet/wrought feedstock, subject to later equivalence checks).

***Emerging Routes*:**

- ***Mid-TRL*:** Multi-step friction stir consolidation (MS-FSC) (solid-state, closed-loop billet-intent). TRL band recorded as mid (≈6–8).
- ***Low-TRL*:** Aluminium hydrolysis to hydrogen (cross-sector, open-loop service displacement; co-produces $Al(OH)_3$). TRL band recorded as low (≈3–5).

For each admitted emerging route, the route register records: (i) TRL band, (ii) intended output form and indicative specification class (e.g., billet-intent; merchant gaseous H_2 intent), and (iii) primary evidence basis (peer-reviewed unit-operation studies and/or audited technical reports). Substitution credits are not applied here. Credits are implemented only after Step 3 defines displaced markets and marginal mixes, and after Step 4 operationalizes equivalence and conditioning in the foreground.

4.2.5 Contemporary Snapshot

A contemporary background snapshot was frozen as the reference for all prospective comparisons. The snapshot corresponded to Brightway project *pCLCA_CA_2025_contemp* with

base database ecoinvent-3.10.1-consequential. No premise transformations were applied. This fixed background ensured that all contemporary inventories and sensitivity results reported here were evaluated against an identical, versioned dataset. Reproducibility materials (manifest and integrity records) were archived in Appendix C.

4.3 Step 2: Scenario-Consistent Backgrounds

This step instantiated the §3.3 rules to produce a set of frozen 2050 prospective background databases for use in subsequent consequential modelling. The contemporary baseline (ecoinvent 3.10.1 consequential; frozen in Step 1) was not reconstructed here. Instead, three IMAGE v3.0 pathways were translated with premise into ecoinvent-compatible consequential databases for the 2050 assessment year.

By design, only the electricity system and fuel/energy-carrier supply system were structurally transformed in the prospective backgrounds. All other background sectors retained their original structural configuration but operated under the transformed energy supply defined by the pathway-specific electricity and fuels markets generated by premise.

To support scenario interpretation under alternative assumptions about anticipatory change, each pathway was instantiated under a premise "foresight" configurations:

- ***PERF (foresight=True):*** anticipatory / perfect-foresight setting used for the primary prospective backgrounds.

This yielded three frozen prospective backgrounds.

4.3.1 Scenario Set & Database Instantiation

Three contrasting IMAGE pathways were selected to span structural uncertainty and were instantiated as separate prospective backgrounds for 2050:

- ***SSP1–VLLO*** *(low forcing; optimistic)*
- ***SSP2–M*** *(medium forcing; central)*
- ***SSP5–H*** *(high forcing; pessimistic)*

Each pathway was built using PERF foresight, producing the following frozen Brightway databases:

- prospective_conseq_IMAGE_SSP1VLLO_2050_PERF
- prospective_conseq_IMAGE_SSP2M_2050_PERF
- prospective_conseq_IMAGE_SSP5H_2050_PERF

4.3.2 Background Construction & Freeze (v1.0 background set)

Each pathway build used ecoinvent 3.10.1 consequential as the base system and premise as the transformation engine, executed inside the Brightway project environment with an ecoinvent-aligned biosphere database available (ecoinvent-3.10.1-biosphere) and LCIA methods present.

For each pathway, background construction followed the same deterministic pipeline, with run metadata written to a manifest and run logs archived (Appendix C):

> ***(i) Temporal Anchoring (single-year snapshot):*** Each pathway background was instantiated as a single-year 2050 operating-state snapshot via premise interpolation. No rolling-window smoothing, averaging across adjacent years, or multi-year blending was performed during background construction. Multi-year IAM series were retained only as upstream evidence and were not treated as multi-year inventories within Step 2.
>
> ***(ii) Scope of Transformation (energy-only structural change):*** Only electricity generation/supply and fuels/energy carriers (e.g., natural gas, refined products, biofuels, hydrogen, synthetic fuels) were updated to their 2050 pathway structure. Other sectors (e.g., cement, steel, transport, industrial heat) were not rewritten at this stage and instead inherit the pathway-consistent energy supply through their electricity/fuel inputs.
>
> ***(iii) Write to Brightway and Preliminary Checks:*** After updating the required sectors, each transformed database was written to Brightway. Pre-freeze checks verified:
>
> - presence of the pathway-specific 2050 electricity and fuel markets generated by premise,
> - biosphere/LCIA resolvability (i.e., biosphere links point to the project biosphere and can be processed),
> - activity counts within expected ranges for a complete prospective database (≈26k activities for the 3.10.1 consequential system under the selected configuration).
>
> ***(iv) Freeze Artifacts:*** Each background was then frozen as part of the v1.0 prospective background set via:
>
> - a manifest capturing configuration, software versions, timestamps, counts, and checksums, and
> - an audit pack of integrity results (Appendix C).

4.3.3 Post-Build Quality Assurance (QA) and "No-Edit" Outcome

Post-build QA was intentionally designed as light-touch integrity verification, not an opportunity for manual restructuring. The objective was to confirm that each premise-generated background could be treated as a frozen reference world for the assessment year.

4.3.3.1 Integrity and Linkage QA (background self-consistency)

Each frozen background was scanned for:

- ***Unexpected External Technosphere Dependencies*** (background should be self-contained),
- ***Broken Technosphere Links*** (inputs that do not exist in the referenced database),
- ***Wrong Biosphere Links*** (biosphere exchanges pointing to a non-project biosphere),
- ***Location Coverage Signals*** (activity location presence/recognition as a proxy for regionalization readiness, while acknowledging premise may still warn about geocollections).

4.3.3.2 Project-Wide Orphan Exchange Prevention

Because Brightway processing and LCA initialization may trigger project-wide database processing, the QA stage also included a project-level scan for orphaned exchanges (exchanges whose input references a database name that does not exist). This check is necessary in workflows where foreground databases are iterated while background database names evolve (e.g., when moving from preliminary builds to final frozen names).

Where orphaned exchanges were detected in foreground databases (i.e., processes pointing to previously deleted or renamed background databases), these were treated as a hygiene issue to be resolved outside Step 2. The corrective action was to relink foreground exchanges to the final frozen background database names (preferred), rather than implicitly treating deletions as an acceptable steady-state solution.

4.3.3.3 Geographic Anchoring: Bridges Not Required

A targeted anchoring check confirmed that activities already resolved to valid pathway-region electricity and fuel markets produced by premise (e.g., IMAGE regional markets such as NA where applicable). As a result:

- no pass-through bridge markets were created, and
- no manual reweighting was performed.
  This was recorded as a "no-action—aligned" QA outcome (Appendix C).

4.3.3.4 Versioning and Freezing Record

Because QA yielded no inventory edits to the premise-generated backgrounds, each database continues to be cited as part of the v1.0 prospective background set. The manifests and checksums provide the citable freeze record. A "frozen" marker and checksum pointer are also written into Brightway database metadata to record intent.

4.3.4 Limitations

Partial sectoral transformation remained an explicit constraint. Cement, steel, heavy transport, industrial heat variants, and direct air capture were not fully rewritten into their 2050 structures. Instead, these sectors operated under 2050 electricity and fuel markets but retained their contemporary structural configuration. This choice may have understated decarbonization in optimistic scenarios (SSP1-VLLO) and understated intensification where fossil lock-in persisted

in pessimistic scenarios (SSP5-H). These sectors were flagged for sensitivity analysis in future work rather than embedded through partial or hand-edited restructuring.

Regional, not subnational, resolution was adopted in the Step-2 prospective backgrounds. IMAGE and premise emitted energy markets at a North America scale. No Canada-specific bridge markets were created, and no provincial grids or province-specific fuel logistics were introduced inside the frozen background databases. Where subnational specificity was required for foreground operations (notably to align province-resolved marginal supplier constructs for hydrogen and primary aluminium), provincial electricity mixes were introduced explicitly in Step 3 as foreground constructs, without altering the frozen Step-2 backgrounds.

Parameter versus scenario uncertainty was separated. The three frozen backgrounds bracketed structural uncertainty across SSP1-VLLO, SSP2-M, and SSP5-H. Within-scenario parameter uncertainty (e.g., yields, loss rates) and decision uncertainty (e.g., alternative EoL routing assumptions) were handled in later steps through foreground modelling and sensitivity analysis. This separation avoided co-evolving the background during consequential evaluation and kept attribution transparent.

4.4 Step 3: Defining Market Boundaries & Computing Marginal Supplier Mixes

This step established, for each retained hotspot material and EoL (reuse, landfill, conventional recycling, MS-FSC, and aluminium hydrolysis), the credited displaced product market and the scenario- and year-specific marginal supplier mix for that market. No credits were assigned here. The output of this step is the marginal supplier vector used in Stage D substitution (§4.5) to translate recovered outputs into avoided production.

In this case study, explicit marginal-mix construction was applied only where required to represent the Canadian decision context for retained hotspot markets and credited products. For the contemporary case, this included: (i) electricity, derived from forward-looking Canadian generation projections (2025–2030) using net-change normalization; (ii) primary aluminium, derived using an empirical “who expands?” screen based on realized output changes (2018–2023); and (iii) industrial gaseous hydrogen, constructed for the hydrolysis displaced-product claim and resolved in the final central case to the single supplier line classified as modern under the net-trend screen. For reuse, the displaced market remained a façade-component/service replacement claim, but because a directly observable reuse-market marginal time series was not available, the final implementation used the same primary-aluminium marginal mix as conventional recycling together with an avoided-extrusion proxy at Stage D. All other materials used ecoinvent consequential markets and supplier structures as provided in the ecoinvent/premise background, without additional foreground marginal-mix construction.

For 2050, electricity and fuels (and hydrogen where represented) were taken directly from the premise-generated IMAGE 2050 PERF scenario databases. No additional retrospective trend detection was applied to these commodity mixes in the case study. For primary aluminium in

2050, plant-level forecasts were not imposed. Instead, the contemporary supplier set was retained as a "who supplies?" construct and a transparent technology-evolution overlay (conventional vs inert-anode/ELYSIS-type) was applied to represent expected change in the production mechanism under 2050 conditions, while upstream energy and trade/import structures were inherited from the premise-transformed background for the selected scenario-year.

All market boundaries and marginal-mix inputs were compiled from data published up to and including 2024 to avoid confounding from post-2024 policy and trade shifts. Any deviation from this rule is documented and justified. Where marginal eligibility relied on heuristic screening (primary aluminium and contemporary hydrogen), the implications were bounded through sensitivity variants (e.g., import-enabled and/or CRR-screened alternatives), reported in Appendix D.

4.4.1 Step 3-1: Scale and Time Horizon of Change

Trend detection relied on fixed, short windows tailored to each context:

- ***Contemporary Window***
    - ***Aluminium (heuristic):*** A fixed 2018–2023 window was adopted as a documented exception for data availability and signal stability.
    - ***Contemporary Hydrogen (heuristic):*** A multi-year realized-output window was applied where available (2020–2024 where ≥3 consecutive observations exist; otherwise, 2024–2025), with all deviations logged.
    - ***Contemporary Electricity (forecast-based):*** A projection horizon of 2025–2030 was used to identify what is expected to respond under long-run planning conditions.
- ***Prospective Window***
    - ***Electricity and Fuels (including hydrogen):*** Marginal mixes were taken directly from the premise-generated IMAGE 2050 PERF scenario database, using premise's embedded consequential identification settings (range time = 2 years, perfect foresight (PERF), CRR off, slope-based method). Accordingly, no additional retrospective trend-detection window was applied in the case study for these commodity mixes.
    - ***Aluminium:*** Plant-level forecasts were not imposed. Contemporary supplier shares (2018–2023 "who expands?" screen) were retained for "who supplies?", while a conventional vs inert-anode (ELYSIS-type) technology overlay represented prospective change. Upstream energy and trade/import structures were inherited from the premise-transformed background for 2050. Electricity and fuels marginal mixes were taken directly from the premise-generated IMAGE scenario database for the selected scenario-year (2050),

These horizons were applied only where the case study required explicit foreground screening or construction (aluminium and hydrogen heuristics, and the electricity projection basis). No component-specific windows or bespoke scaling were introduced beyond the documented exceptions, and all deviations were recorded in the audit log.

4.4.2 Step 3-2: Limits of a Market

This subsection fixed the displaced product market for each end-of-life (EoL) route by stating the largest defensible market the recovered output could serve. It established only the boundary conditions: (i) functional equivalence, (ii) grade/specification where applicable, and (iii) consequential geography and trade routing. It did not specify conditioning/upgrading operations, yields, QA evidence, or quality discounts. Those elements were handled in foreground EoL modelling and Stage D substitution.

As landfilling created no credited market, four candidate crediting routes remained for aluminium processing at EoL:

1. ***Reuse*** (in the same capacity) — *Displaced Market*: Aluminium non-structural façade components.
2. ***Conventional Recycling*** (Closed-Loop) — *Displaced Market*: Wrought extrusion feedstock (billet).
3. ***Mid-TRL Emergent Recycling*** (Closed-Loop) — *Displaced Market*: Wrought extrusion feedstock (billet).
4. ***Low-TRL Emergent Recycling*** (Open-Loop) — *Displaced Market*: Industrial gaseous merchant hydrogen.

Market limits followed the rules in § 3.4. Where residual grade/spec uncertainty was plausible. A condensed justification for each market follows including full boundary notes, geography rationales, and summation of this definition and relevant sources are provided in Appendix D.

4.4.2.1 Reuse

The displaced market was aluminium non-structural façade components where fit, geometry, and service duty matched the recovered items, consistent with CLCA boundary heuristics that prioritize functional equivalence at the market-definition stage. Alloy chemistry did not determine eligibility at this step because reuse “as-is” turns on the obligatory properties of the displaced product rather than metallurgical composition.

Consequential geography was set at the Canada-wide level. Although reuse flows are often organized locally through deconstruction, diversion, and project-specific redistribution channels, cross-provincial movement and broader distribution networks make a national boundary the most defensible scale for identifying marginal supply, while local logistics are handled separately in the foreground. Canadian policy and sector reviews document reuse as locally executed but embedded within a broader materials system, which supports defining the displaced market at the national level where suppliers compete and where trend evidence can be observed.

In the final implementation, this market boundary remained a façade-component/service claim rather than being redefined as a billet market. However, because a directly observable Canada-wide marginal time series for reusable façade-component supply was not available, the practical

marginal-response representation used for Stage D was bounded using the same primary-aluminium marginal mix adopted for conventional recycling, paired with the avoided-extrusion proxy. Similar-product imports were screened and were not retained as marginal contributors because they represented less than 0.1% of total production. The final reuse claim should therefore be interpreted as a service-market claim implemented through the closest defensible observable aluminium supply proxies, rather than as a direct observation of a national reuse-market marginal mix.

4.4.2.2 Conventional Recycling

The displaced market was wrought extrusion feedstock in billet form supplied to 6000-series extrusion. Functional equivalence was defined at the product level as extrusion billet conforming to recognized 6000-series alloy limits (AA/EN) and suitable for extrusion service after standard cast house conditioning (The Aluminium Association, 2023; CEN, 2019). This definition is agnostic to metal origin (primary vs recycled) and focuses on the billet specification itself. Additionally, this boundary is supported by industry practice and technical literature showing that extrusion cast houses routinely produce billets from process and recycled scrap and that, with appropriate scrap selection and melt refining, secondary (recycled) billets can meet 6000-series compositional and quality targets required for extrusion markets (Al-Alimi et al., 2024).

Consequential geography for the market limit is set at the Canada-wide level. This choice reflects the scale at which billet suppliers compete, as opposed to primary ingot trade, and the ease of cross-provincial movements that are not consistently tracked at the level needed to construct reliable provincial marginal mixes. Import origins and supplier composition are resolved later through exporter-bin routing in Step 3-3/3-4. These downstream results do not alter the market boundary fixed here.

4.4.2.3 Mid-TRL Emergent Recycling

The displaced product was wrought extrusion feedstock in billet form suitable for 6000-series extrusion, consistent with the conventional recycling case. Because commercial billet is transacted by obligatory properties specified in alloy and temper standards rather than by "structural versus non-structural" end-uses at the point of sale, no application split was imposed at Step 3-2 (Consequential-LCA, 2015; Weidema et al., 2004). Consequential geography followed the Canada-wide rule for billet supply, consistent with defining the market where marginal suppliers actually compete.

4.4.2.4 Low-TRL Emergent Recycling

The displaced product is merchant gaseous hydrogen sold for industrial use. Functional equivalence is defined at the product level as hydrogen delivered to customers as a bulk industrial gas, independent of production route. In Canada, current hydrogen demand is dominated by refining and chemical uses, and merchant hydrogen is commonly supplied as gaseous product rather than defined by a specific end-use label. This product-level framing

aligns with C LCA guidance to define the market by the goods that meet marginal demand rather than by the process history (Ekvall & Weidema, 2004).

Consequential geography is set at the Canada-wide level. Industrial hydrogen supply competes at national scale through a combination of on-site production, regional pipeline networks in established hubs (e.g., Sarnia and Alberta's Industrial Heartland), and tube-trailer distribution. These logistics and infrastructure allow suppliers across provinces to serve marginal demand where it arises, whereas provincial-only marginal mixes cannot be robustly constructed due to cross-provincial flows and sparse, inconsistent tracking. Pipeline hubs and merchant distribution are well documented, and detailed logistics and distances are carried forward to Step 4 (NRCan, 2024; US DOE, 2025).

Application splits are not imposed at Step 3-2 because industrial hydrogen is not transacted by application category at the point of sale. Claims are therefore framed against the gaseous hydrogen market. Any end-use-specific purity or pressure adjustments are handled in the foreground and substitution logic in later steps. Aluminium–water hydrolysis yields hydrogen and co-produces aluminium hydroxide. The latter was mapped separately at Stage D to the appropriate market with geography consistent with the consequential boundary with further details found in Appendix D.

These market limits applied unchanged in contemporary (2025) and prospective (2050) analyses; they define which market is displaced, not how equivalence is achieved. Prospective differences arise through the frozen Step 2 backgrounds (evolving efficiencies and electricity/fuel mixes) and through the foreground marginal supplier composition derived in Step 3-3/3-4. Any conditioning/upgrading required to meet the stated market specification, and any quality-correction or remapping where evidence is partial, was handled in Step 4 and Stage D substitution (§ 4.5).

4.4.3 Step 3-3: Trends in the Volume of a Market

This subsection establishes whether observed volume changes over the detection window are large enough, relative to capital replacement, to identify which suppliers are expanding capacity in practice. The analysis is deliberately limited to volume trends as it does not assign credits, model substitutions, or infer consequential demand shifts. Those elements follow in Step 3-4.

As established in Step 3-2, several displaced markets arise from the different EoL pathways of the aluminium components within the functional unit. In what follows, we treat each displaced market in turn, beginning with recycled billet, then re-using the evidence base to motivate reuse (where feasible), and closing with a brief scan of hydrogen-related trends. Prospective trend detection for the aluminium EoL pathways is reported in Appendix D.

4.4.3.1 Market Trends for Contemporary Aluminium Billets

For trend detection we defined the relevant product market as unwrought, extrusion-suitable aluminium (e.g., 6000-series billet precursor). The supplier universe was compiled as: (i) domestic primary smelters, indexed by $i \in D$, resolved at facility level; and (ii) import partners, indexed by $j \in M$, resolved bilaterally and filtered to billet-equivalent unwrought aluminium. Here, D denotes the set of domestic smelters (facility-level bins) and M denotes the set of import partners (country-level bins). Secondary/remelt capacity was treated as constrained at this step and therefore excluded from long-term capacity-margin detection, while retained for short-term operational sensitivities reported later. This treatment reflects CLCA guidance that constrained technologies do not expand capacity to meet a sustained demand change and therefore are not long-run marginal suppliers. Secondary/remelt production was treated as supply-constrained by scrap availability and quality, consistent with CLCA guidance that scrap-based routes do not expand to meet sustained demand changes. Accordingly, secondary production was excluded from long-run capacity-margin detection and retained only for short-term operational sensitivities (Beylot, 2024).

Because plant-level capacity additions and investment signals were not available consistently by facility, long-run marginal eligibility was proxied using multi-year changes in realized output as an indicator of which facilities were operationally positioned to supply additional volume. This is treated as an empirical screening heuristic.

To maintain transparency, the classification was executed in two passes using the same detection window and metric: first on domestic production only, then repeated after admitting imports that meet the billet-equivalent filter. This sequencing isolates the incremental effect of imports on the composition of expanding capacity, consistent with consequential market logic in which the long-run marginal suppliers are those that invest (or disinvest) in response to the market-wide change.

Domestic Production

In Canada, aluminium is not mined as bauxite as in the typical aluminium process. Instead, alumina is imported and refined, then smelted to primary aluminium most commonly via the Hall-Héroult process (Aluminium Association of Canada, 2025). Approximately four to five tonnes of bauxite yield about two tonnes of alumina, and roughly two tonnes of alumina yield one tonne of primary aluminium. In 2023, Canada operated ten primary smelters (nine in Québec, one in Kitimat, BC) and one alumina refinery (Jonquière, QC). These 10 smelters generated a national output of primary aluminium of ~3.3 million tonnes in 2023 (NRCan, 2024). This matched an additional facility-by-facility tally of roughly 3.275 million tonnes. These facts support a facility-resolved approach to trend detection and, in later steps, a province-resolved treatment of electricity inputs because power mixes and their trajectories differ by location (Government of Canada, 2023).

To compute the domestic supplier universe, a 2018–2023 window was used to classify which suppliers exhibited sustained capacity increase large enough, relative to normal replacement, to count as expanding in practice. For each facility, the annualised change between the window endpoints was computed and compared with a replacement load derived from the smelter's technology lifetime, which is typically measured in multiple decades. Within the window, a structural break occurred in Rio Tinto's Aluminerie Arvida (formerly know as Chicoutimi) complex due to an expansion (Arvida AP60) alongside the gradual closure of pot rooms at the legacy Arvida smelter (Rio Tinto, 2023). This renovation-driven curtailment was flagged so that it would not be misread as a competitiveness-driven contraction. In addition, the Rio Tinto Beauharnois facility has ceased operation since the end of the recorded window. Post-window events were treated as scenario-year eligibility constraints rather than retroactive changes to detection. The detection metric was applied facility-by-facility. The corresponding plant list, volumes, and calculations are provided in Table 4-2 with the formal expression in Equation 4-1.

***Table 4-2**: Primary aluminium (billet-equivalent proxy) – window volumes and expansion signal by historic domestic producer (2018-2023)*

Supplier	***Location***	***Technology***	***Output (kt) 2018***	***Output (kt) 2023***	***ΔV (kt)***	***Annual ΔV (kt/yr)***
Rio Tinto – Alma Smelter	Quebec	ELYSIS/AP 44	465	484	19	3.8
Aluminerie Alouette – Alouette Smelter	Quebec	Prebake	584	634	50	10.0
Alcoa – Deschambault Smelter	Quebec	Prebake	260	287	27	5.4
Rio Tinto – Grande Baie Smelter	Quebec	Prebake	233	229	-4	-0.8
Alcoa – Becancour Smelter	Quebec	Prebake	438	465	27	5.4
Rio Tinto – Kitimat Smelter	British Columbia	AP40	436	377	-59	-11.8
Total			**1,980**	**2,476**	**60**	**12.0**

Equation 4-1

$$\Delta V_i = V_i^{end} - V_i^{start},\ i \in D$$

Where V_i^{start} and V_i^{end} are the first/last annual volumes in the detection window.

Facilities were treated as "included in the marginal set" when $\Delta V_i > 0$. This mix represents marginal expansion signals over 2018–2023, not total production shares. The domestic pass classified facilities with $\Delta V_i > 0$ as included in the marginal set and recorded technology class and province for downstream use in Step 3-4.

Imports

Having established that the domestic set did not expand faster than replacement, the same detection was repeated after admitting imports. In CLCA, the marginal mix is derived for the product market that meets the change in demand under the stated boundary. When local capacity is not expanding sufficiently, imports belong to the long-run margin when they are functionally equivalent to the displaced product and show a positive trend within the detection window. This placement follows the "increasing or slowly decreasing market" rule and the aluminium example, which bases marginality on who expands capacity rather than nationality.

Canada monitors specified aluminium products under General Import Permit No. 83 through the Aluminium Import Monitoring Program, which provides partner-by-period datasets (ALU1, ALU2) (Global Affairs Canada, 2020). For a billet-equivalent unwrought market, the screen retained HS 7601 at the HS-10 level (Appendix D), including ingot and "other" statistical lines that, in Canadian reporting, encompass billets and slabs. Semi-fabricated forms were excluded (bars, rods, profiles, wire, plate, sheet, strip, tube, fittings; HS 7604–7609). The monitoring program's HS list and permit materials were used to document the filter and its rationale.

Import bins were defined by partner country. The identical window and expansion-signal metric used for domestic facilities was applied to each partner-country series after the HS filter and routine conditioning (e.g., very small or sporadic flows handled per Appendix D), illustrated in Equation 2. This expansion-signal screen was used to test whether any import partners contributed materially to marginal expansion over 2018–2023. In this case, qualifying import growth was non-zero but negligible relative to domestic expansion signals. Import partners were therefore retained as an explicit check but did not materially influence the contemporary marginal mix (full results in Appendix D). Accordingly, the central case marginal supplier set for Step 3-4 was dominated by domestic facilities with $\Delta V_i > 0$, with imports contributing only a minor residual where applicable. Location continues to matter in the next step for conditioning foreground energy and logistics, while the displacement geography remains the product-market level. For 2023, Canada's aluminium imports were valued at about 9.5 billion dollars, underscoring the quantitative role of external supply in gross trade value, even where screened billet-equivalent import expansion, Equation 4-2, was small over the detection window.

Equation 4-2

$$\Delta V_j = V_j^{end} - V_j^{start},\ j \in M$$

For 2050, aluminium prospective modelling used the Canada-anchored technology-evolution overlay approach (§3.5). In this formulation, the 2025 Canada + importer-bin supplier set was retained as the supplier-identity structure, while prospective change was represented through long-horizon adjustments to supply emissions intensity, rather than through projected changes in facility-level market shares. Because credible plant-level forecasts for billet-equivalent imports were unavailable, the central prospective case retained imports at their last-observed negligible contribution rather than imposing speculative growth.

4.4.3.2 Market Trends for Contemporary Extruded-Aluminium

Extruded-aluminium evidence was compiled as supporting evidence for the plausibility of the reuse claim and for the avoided-extrusion component of the Stage D proxy. However, the final deterministic implementation did not retain a separate standalone reuse marginal mix based on extrusion producers. Instead, reuse was implemented as a façade-component/service replacement claim represented through the combined ingot-plus-extrusion proxy, with the primary-aluminium marginal mix aligned to the same Canada-wide marginal representation used for conventional recycling. Similar-product imports were screened and found to be less than 0.1% of total production and were therefore not retained as marginal contributors in the final central case.

4.4.3.3 Market Trends for Contemporary Industrial Gaseous Hydrogen

Supplier lines for industrial gaseous hydrogen were first defined at the level of identifiable Canadian producer nodes. These included natural gas reforming and steam methane reforming (SMR) facilities, electrolytic hydrogen projects, and methane-pyrolysis developments where sufficient public information was available. Where possible, projects were tagged by province to preserve linkage to the underlying electricity and natural-gas systems. Importer bins were not resolved explicitly because publicly accessible federal statistics and regulatory sources did not identify hydrogen imports at a level that could be mapped consistently to specific foreign producer regions or facilities (Government of Canada, 2023; NRCan, 2024).

In practice, public information on Canadian hydrogen production is sparse, heterogeneous, and unevenly resolved across suppliers. The supplier list used here was assembled from NRCan's Hydrogen Strategy for Canada and its 2024 progress reporting, provincial hydrogen strategies, company disclosures, regulatory summaries, and industry compilations. These sources explicitly note that project information is assembled from publicly available material, proponent submissions, and commercial data sources (NRCan, 2024). As a result, plant-level capacities and annual production volumes for 2020–2025 should be interpreted as indicative rather than exhaustive.

Rather than collapsing suppliers immediately into broad technology families, the final screening retained the candidate supplier lines at the producer level wherever possible so that the net-trend test could be applied transparently to the assembled evidence base. The final candidate set comprised Alberta Energy Regulator natural gas reforming, Alberta Energy Regulator SMR, Air

Liquide electrolysis (Bécancour, QC), HTEC electrolysis (Burnaby, BC), Aurora Hydrogen methane pyrolysis (Fort Saskatchewan, AB), Canadian Natural Resources Limited SMR (Wood Buffalo, AB), and Messer Canada electrolysis (Magog, QC).

Due to data limitations, 2025 was additionally used as a reference year to capture the direction of change for suppliers with short historical series. For Alberta production, multiple years of reported output were available, allowing a longer trend basis. For several non-Alberta and emerging suppliers, however, production or capacity observations were only available from 2024 onward. These shorter records were therefore treated as evidence of presence and directional emergence, but not as equally strong evidence of sustained marginal response.

The resulting volumes were benchmarked against national estimates as a plausibility check. NRCan reports Canadian hydrogen production on the order of 4 million tonnes per year in 2024, with Alberta representing the dominant share. The assembled supplier-line totals used here are therefore best interpreted as a bounded screening dataset for marginal inference rather than as a full national production census.

For each supplier line, the trend metric in Equation 4-1 was then applied over the available window, using 2020–2025 where sufficient observations existed and shorter windows only where earlier data were unavailable. Net annual trend was evaluated after applying the capital-replacement correction defined in Step 3-4. Under this screen, only the Alberta Energy Regulator steam methane reforming (SMR) plant was classified as a modern technology and therefore was solely retained in the contemporary marginal set. All remaining candidate lines, including Alberta natural gas reforming, Air Liquide electrolysis, HTEC electrolysis, Aurora methane pyrolysis, Canadian Natural Resources Limited SMR, and Messer Canada electrolysis, were classified as old and received zero marginal share in the final central case.

Accordingly, the contemporary hydrogen marginal mix used for the hydrolysis displaced-product claim was 100% Alberta Energy Regulator SMR. This should be interpreted as a conservative central-case marginal result under the available evidence, not as a claim that other hydrogen pathways are irrelevant in Canada. Rather, it reflects that only one supplier line exhibited a positive net annual trend large enough, and with sufficient temporal support, to qualify as modern under the implemented screening rule.

4.4.4 Step 3-4: Changes in Supply and Demand

This section converted the Step 3-3 trend signals into a scenario-year marginal supplier mix for each hotspot market defined in Step 3-2. The joint supplier set was $S = D \cup M$, combining domestic bins $D$and importer bins M. For each supplier $i \in S$, the annualized fractional growth rate over the detection window was

Equation 4-3

$$g_i = \frac{V_{i,t_1} - V_{i,t_0}}{V_{i,t_0}\,\Delta T}$$

with V_{i,t_0} and V_{i,t_1} denoting harmonized volumes at the window start and end, and ΔT the window length in years. This definition applied equally to domestic and importer bins once volumes were harmonized to the market-availability basis stated in Step 3-3.

To distinguish genuine capacity expansion from simple renewal of aging assets, a capital-replacement requirement was computed using a technology lifetime, L_i, appropriate to supplier i's dominant equipment:

Equation 4-4

$$r_i = -\frac{1}{L_i}$$

For domestic suppliers ($i \in D$), $L_i$was resolved at facility or technology-class level. For importer bins ($j \in M$), exporter-specific lifetimes were not resolved, so $r_j$was set to 0 by design (imports carry no replacement correction).

The net annual trend that governed marginal eligibility was

Equation 4-5

$$\tau_i = g_i + r_i$$

Positive τ_i indicated net expansion above replacement within the window, whereas negative τ_i indicated capacity at or below replacement. Where documented structural breaks occurred within the window, series were split and $\tau_i$was computed segment-wise to avoid misclassifying renovation-driven dips as loss of competitiveness. Lifetime values L_i, data sources, imputation rules, and all breakpoints were logged in Appendix D. For interpretability across markets, technologies with $\tau_i > 0$ were labelled modern (adding net effective capacity) and those with $\tau_i \leq 0$ were old within the window. This visual classification is referenced in Figure 4-2.

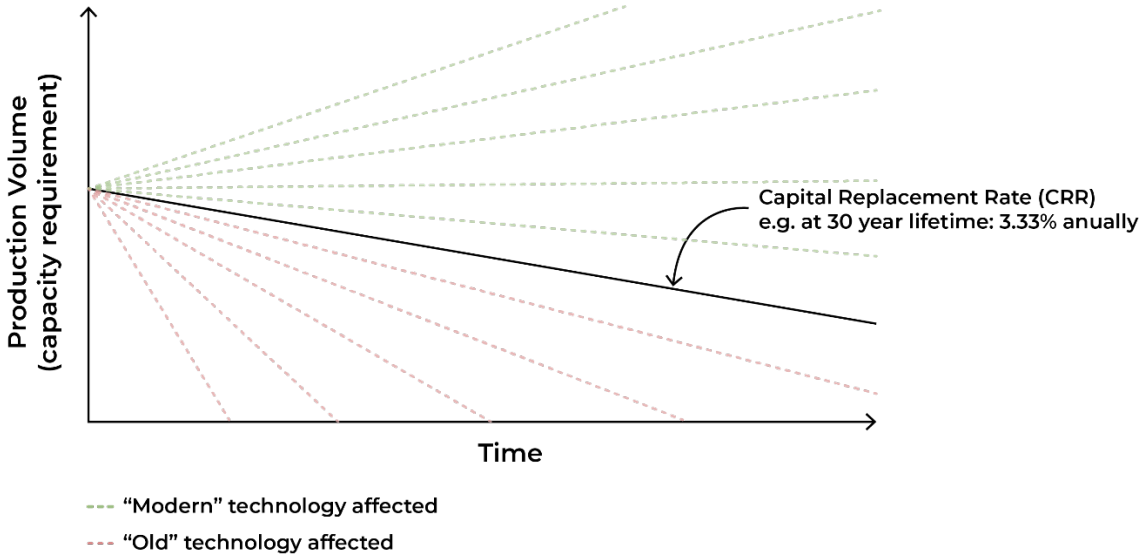


***Figure 4-2**: Conceptual illustration of the capital replacement rate (CRR) used to distinguish suppliers that are net expanding (green dashed trajectories; "modern" technologies contributing to marginal supply) from those that are net contracting (red dashed trajectories; "old" technologies). The black line represents the CRR threshold (e.g., 3.33% per year for a 30-year lifetime), showing the minimum growth needed to exceed replacement and indicate effective capacity expansion over time. Derived from (Consequential-LCA, 2020).*

For aggregation at market level, the net trend $\tau_i$was treated as a fractional rate, and the corresponding absolute net change for supplier $i$was obtained by scaling with the window-start volume:

Equation 4-6

$$\Delta C_i \;=\; \tau_i \, V_{i,t_0}$$

The market-level signal was then the sum of absolute net changes $\sum_{i\in S} \Delta C_i$. If $\sum_i \Delta C_i > 0$, the market was expanding above replacement, and the co-marginal set was the modern suppliers $\mathcal{M} = \{\, i \in S : \Delta C_i > 0 \,\}$. Shares were allocated in proportion to absolute net expansion and normalized,

Equation 4-7

$$\alpha_i \;=\; \frac{\Delta C_i}{\sum_{k\in\mathcal{M}} \Delta C_k} \;(i \in \mathcal{M})$$

If $\sum_i \Delta C_i \le 0$, the window was contracting and the co-marginal set consisted of the least-declining (old) suppliers; shares were apportioned from the same source-family baseline composition but restricted to that set, then normalized (Equation 7). Before normalization, post-

window feasibility updates for the scenario year such as closures, sustained curtailments, or commissioning were applied to the relevant modern/old bins with justification.

This procedure yielded the scenario-year marginal mix $\{\alpha_i\}_{i\in\mathcal{M}}$ for the market. Market-specific results, tables, and figures are provided in the following subsections, and the full calculation trail $(V_{i,t_0}, V_{i,t_1}, \Delta T, L_i, g_i, r_i, \tau_i, \Delta C_i)$, together with structural-break logs and feasibility updates, was consolidated in Appendix D.

Importer bins used the observed bilateral growth as the net trend, $\tau_j = g_j$ and $\Delta C_j = g_j V_{j,t_0}$. This avoided unsupported inference about foreign capital stocks while keeping aggregation internally consistent.

Equation 4-1 to Equation 4-7 construction and data is reported for the contemporary (historically inferred) marginal mixes. Prospective electricity and fuels marginal mixes are inherited directly from the premise/IAM scenario database for the selected year and are therefore not recomputed using the historical screening equations in this section. Prospective primary aluminium retains the contemporary "who supplies?" structure, with a technology overlay applied to represent inert-anode/ELYSIS-type evolution under 2050 conditions, documented fully in Appendix D.

4.4.4.1 Market Trends for Contemporary Aluminium Billets

The resulting contemporary marginal mix for aluminium billets follows in Table 4-3.

***Table 4-3**: Contemporary marginal supplier mix for the aluminium billet market (Canada; 2018–2023 expansion-signal basis)*

Supplier	***Annualized growth, (g_i) (%/yr)****	***Positive Annual ΔV, ΔV^+ (kt/yr)***	***Included in marginal set?***	***Marginal mix (%)***
Rio Tinto – Alma Smelter	0.82%	3.80	Yes	15.45%
Aluminerie Alouette – Alouette Smelter	1.71%	10.00	Yes	40.65%
Alcoa – Deschambault Smelter	2.08%	5.40	Yes	21.95%
Rio Tinto – Grande Baie Smelter	-0.34%	0.00	No	0.00%
Alcoa – Becancour Smelter	1.23%	5.40	Yes	21.95%
Rio Tinto – Kitimat Smelter	-2.71%	0.00	No	0.00%
Total (ΔV^+ normalized)		**24.60**		**100.00%**

Given that the suppliers included in the contemporary aluminium billet marginal mix are concentrated overwhelmingly in Quebec, province-specific electricity inputs were used to represent foreground electricity demand, rather than applying a national average electricity mix. This geographic conditioning reflects the fact that most Canadian primary aluminium manufacturing occurs in Quebec and is consistent with the electricity special case treatment described in § 4.4.4.4. The supporting provincial marginal electricity calculations are provided in Appendix G.

4.4.4.2 Market Trends for Contemporary Extruded-Aluminium

Reuse was assumed to be locally executed, but the marginal supply that could meet like-for-like profile demand was proxied by domestic extrusion producers and qualifying extrusion imports screened in Step 3-2/3-3. The same τ_i logic and feasibility-before-normalization rule were applied, results are reported in Appendix D.

4.4.4.3 Market Trends for Contemporary Industrial Gaseous Hydrogen

The expanding bins identified in §4.4.3.3 were carried directly into the contemporary hydrogen marginal-mix calculation rather than being re-aggregated into broad technology families for the final central case. This is important because the implemented net-trend screen was applied at the producer-line level and the resulting marginal inference is therefore line specific rather than family based.

For each supplier line, annualized growth, capital-replacement correction, net annual trend, and absolute net annual change were calculated using the same logic applied elsewhere in Step 3-4. Under this producer-line screen, only the Alberta Energy Regulator SMR line exhibited positive net expansion above replacement and was therefore classified as modern. Its observed supply change, after correction for capital replacement, yielded the only positive net annual change among the candidate lines and thus the only non-zero marginal share in the central case. All other candidate lines were classified as old because their observed growth was zero, negative, or insufficient to exceed replacement requirements under the available time series.

Accordingly, the contemporary hydrogen marginal mix was represented as 100% Alberta Energy Regulator SMR in the final implementation with the full calculation trail $(V_{i,t_0}, V_{i,t_1}, \Delta T, L_i, g_i, r_i, \tau_i, \Delta C_i)$ and the excluded low-share families documented in Appendix D. Because the sole included marginal supplier line was Alberta based, the associated contemporary energy inputs were likewise aligned to Alberta-specific conditions rather than a Canada-wide average, preserving consistency between the marginal supplier geography and the underlying production-energy system. The hydrogen central case used for the hydrolysis displaced-product claim should therefore be read as a producer-line-specific marginal result under constrained evidence, rather than as a broader technology-family result.

4.4.4.4 Contemporary Electricity Supply

Electricity was not a credited product in this study, but it materially affected foreground burdens for contemporary and prospective EoL operations. Because electricity is a strategic product regulated largely through national energy policy and planning, the market boundary for marginal electricity was defined at the Canada-wide level. A Canada-wide marginal mix was computed as the default electricity supply for foreground operations, and province-specific marginal mixes were computed only where other marginal supplier constructs (notably hydrogen and primary aluminium) were resolved at the provincial level and required consistent provincial electricity alignment.

For the contemporary case, supplier bins were defined as Canada–technology pairs (for example, Canada–reservoir hydro, Canada–natural gas combined cycle, Canada–nuclear). Consistent with the electricity "special case" guidance that political plans and forecasts are often more relevant than historical competitiveness, the contemporary marginal electricity mix was derived from forward-looking Canadian projections over a 2025–2030 planning horizon (rather than retrospective 2018–2024 trends). Because these series already reflect planned expansion within the projection horizon, no capital replacement rate (CRR) correction was applied in this projection-based construction.

For each technology bin i, projected generation volumes $V_{i,2025}$ and $V_{i,2030}$ were used to compute a net change $\Delta V_i = V_{i,2030} - V_{i,2025}$. The market-level signal was $\Delta V_{\text{mkt}} = \sum_i \Delta V_i$. When $\Delta V_{\text{mkt}} > 0$, the marginal set consisted of technologies with $\Delta V_i > 0$ and shares were allocated by normalizing positive net changes, $\alpha_i = \Delta V_i / \sum_{k:\Delta V_k > 0} \Delta V_k$. When $\Delta V_{\text{mkt}} \leq 0$, the marginal set consisted of technologies with $\Delta V_i < 0$ and shares were normalized analogously over the contracting set.

U.S. electricity imports were not included as a separate marginal contributor in the contemporary projection-based mix because the projection series used did not provide a changing import term over 2025–2030 (that is, imports were effectively treated as stagnant with $\Delta V_{\text{imports}} \approx 0$), and therefore did not contribute to marginal expansion in the net-change normalization.

Where province-specific electricity was required (notably to align province-resolved marginal supplier constructs for hydrogen and primary aluminium), the same projection-based net-change normalization procedure was repeated using the corresponding provincial projections and technology bins. Outside those cases, the Canada-wide marginal electricity mix was applied as the default to preserve consistency across foreground operations. The resulting contemporary Canada-wide marginal electricity mix is reported in Table 4-4 as technology-level shares. The full calculation trail and any constraint assumptions by technology class are documented in Appendix D along with the provincial electricity variations.

***Table 4-4**: Canada-wide marginal electricity mix derived from 2025–2030 projected generation changes (technology-level shares)*

Supplier	GWh (2025)	GWh (2030)	ΔV (2030–2025), GWh	ΔV^{+} = max(ΔV,0), GWh	Marginal share (%)
Hydro / Wave / Tidal	403,350.84	437,850.35	34,499.51	34,499.51	40.09%
Wind	57,101.81	80,824.64	23,722.83	23,722.83	27.56%
Natural Gas	106,529.07	128,015.91	21,486.84	21,486.84	24.97%
Uranium	78,631.37	82,253.63	3,622.26	3,622.26	4.21%
Biomass / Geothermal	7,838.55	9,204.80	1,366.25	1,366.25	1.59%
Solar	12,140.31	13,504.24	1,363.93	1,363.93	1.58%
Coal & Coke	8,132.43	654.81	-7,477.62	0.00	0.00%
Oil	1,379.91	683.79	-696.12	0.00	0.00%
Total	**675,104.29**	**752,992.17**	**77,887.88**	**86,061.62**	**100.00%**

For the 2050 prospective case, the electricity system was already represented in Step 2 by a Canada-level IAM/premise scenario aimed at a deeply decarbonised grid. Because this scenario encodes long-term capacity response consistent with the assumed pathway, separate projection-window trend detection was not re-applied in Chapter 4. Instead, technology shares for Canada in 2050 were taken directly from the selected IAM/premise pathway and mapped to the corresponding Canadian or North American generation processes in the background database.

4.5 Step 4: Model Foreground EoL Pathways and Stage D Substitution

This section instantiated the admitted route set on top of the frozen backgrounds of §4.3 (Step 2) and marginal mixes from §4.4 (Step 3). For each route, it defined the unit-process chain, the equivalence checkpoints, any required proxies, and the Stage D mapping. Case-specific parameters, tables, and worked examples are provided in Appendix B.

4.5.1 Foreground EoL Route Modelling

At deconstruction, aluminium from the curtainwall FU was released and treated as a single “gate scrap” inventory. To represent plausible physical end-of-life states without implying routing shares, the aluminium was characterized using three non-quantitative condition labels that mirror construction practice: (i) intact units potentially suitable for reuse, (ii) clean metallic scrap suitable for recycling, and (iii) contaminated or fine residues that would typically require disposal or additional preprocessing. These labels were used only to enforce route prerequisites (e.g., reuse requires intact units) and to keep unit-process wrappers internally consistent; they were not used to prescribe plausibility, frequencies, or partitioning across routes.

Each route was implemented via a standardised route wrapper with a 1 kg gate-scrap reference product, chaining the necessary unit processes and applying Stage D credits only after

equivalence was reached. Wrappers were then scaled to 3.67 kg per m². to reflect the quantity within the representative FU.

4.5.1.1 Route Mechanics

For aluminium in this case study, five mutually exclusive routing pathways were distinguished:

1. Landfill (baseline disposal).
2. Direct reuse of components.
3. Conventional recycling via remelting.
4. Mid-TRL solid-state recycling via multi-step friction stir consolidation (MS-FSC).
5. Low-TRL chemical recycling via aluminium hydrolysis to hydrogen.

Landfill, reuse, and remelting-based recycling reflect mature, commercially deployed practices, whereas MS-FSC is best characterized as demonstration/early industrial scale and hydrolysis as advanced laboratory scale (Ingarao et al., 2023; Urbonavicius et al., 2024). Both emerging routes were modelled as physically plausible extensions of existing collection, sorting, and scrap-preparation infrastructure, i.e., the novel unit operations were embedded downstream of standard scrap handling rather than assuming an entirely new logistics chain.

Landfilling

Landfilling was treated as the default fate when aluminium scrap was not separated or when recycling was not economically viable. In practice this corresponded to mixed C&D waste sent to municipal or engineered landfills with gas capture. Physical operations were limited to collection, transport, and disposal to reflect C3–C4.

Reuse

Reuse occurred when curtainwall units could be directly reinstalled in another building with minimal refurbishment (D. R. Cooper & Allwood, 2012). In reality this would have required careful deconstruction, visual inspection, and limited repairs, which were excluded from the system boundary. The aluminium was not remelted or chemically transformed. The route extended service life through relocation logistics rather than metallurgical processing.

Conventional Recycling

Conventional recycling via remelting was treated as the dominant route for aluminium scrap. Collected profiles and off cuts were sent to scrap yards or specialised processors, where they were cut, shredded and sometimes de-coated to remove paints and sealants. Prepared scrap was then charged to a natural-gas-fired remelting furnace. Oxidation and dross formation caused a small mass loss, and alloy adjustments could be made with primary aluminium or master alloys

to meet billet specifications. This route returned aluminium to the cycle as secondary billet or extrusion feedstock (Boin & Bertram, 2005). Figure 4-3 illustrates the sectoral flow context.

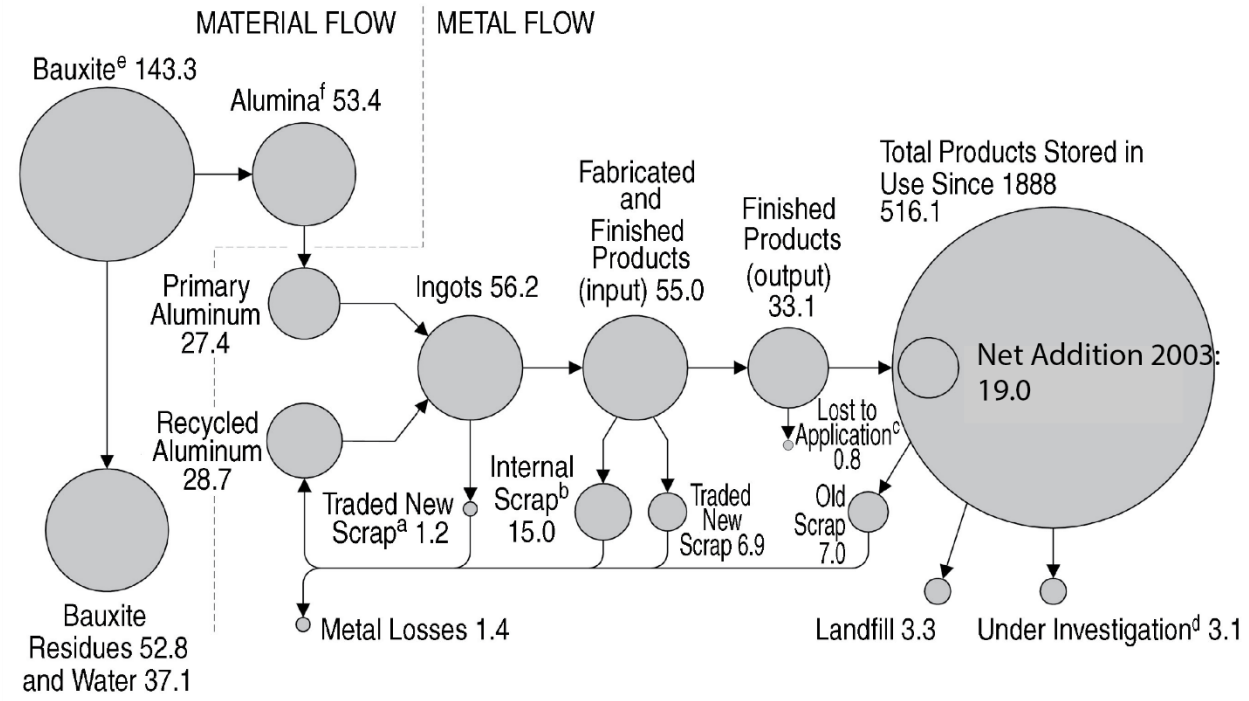


Figure 4-3*: Illustrative Global Aluminium Flow (from Boin & Bertram, 2005). All units measured in million tonnes per year. Values might not add up due to rounding. a–Aluminum in dross; b–not taken into account in statistics; c–such as powder, paste, and deoxidation aluminum (metal property is lost); d–area of current research to identify final aluminum destination (reuse, recycling, or landfilling);e–includes, depending on the ore, between 30% and 50% alumina; f–includes, on a global average, 52% aluminum.*

Emergent Recycling (Mid-TRL): Multi-Step Friction Stir Consolidation (MS-FSC)

MS-FSC was modelled as a solid-state consolidation route that converts cleaned, size-reduced aluminium scrap into a near-dense billet without remelting. In the modelled chain, prepared scrap was reduced to chips/fines, packed into a die cavity, and consolidated through multiple tool passes under confinement. The multi-step configuration was selected (rather than a single-pass variant) because it is commonly reported as improving bonding continuity and billet homogeneity, albeit at higher electricity demand (Ingarao et al., 2023). The consolidated billet was then ejected, trimmed, and minimally finished as an extrusion-grade billet precursor. Figure 4-4 depicts this sequence.

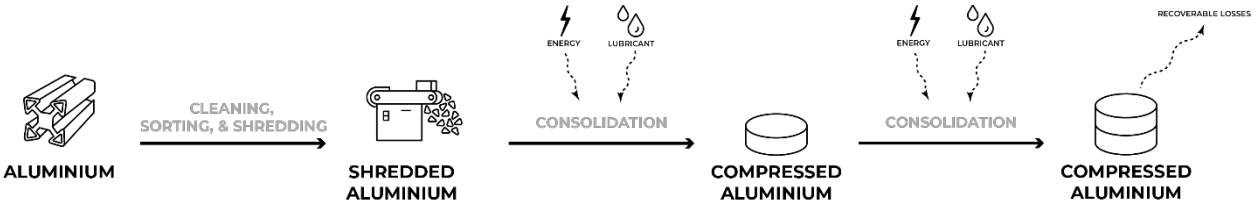


Figure 4-4: Illustrative Single & Multi-Step Friction Stir Consolidation Process (derived from Ingarao et al., 2023)

Emergent Recycling (Low-TRL): Aluminium Hydrolysis

Aluminium hydrolysis was modelled as a low-TRL conversion route in which aluminium reacts in alkaline water to generate hydrogen gas and aluminium hydroxide. The modelled chain comprised scrap preparation, alkali solution make-up with deionized water, exothermic reaction, and recovery of aluminium hydroxide from the alkaline liquor (Urbonavicius et al., 2024). Figure 4-5 outlines the input–output structure and process interfaces.

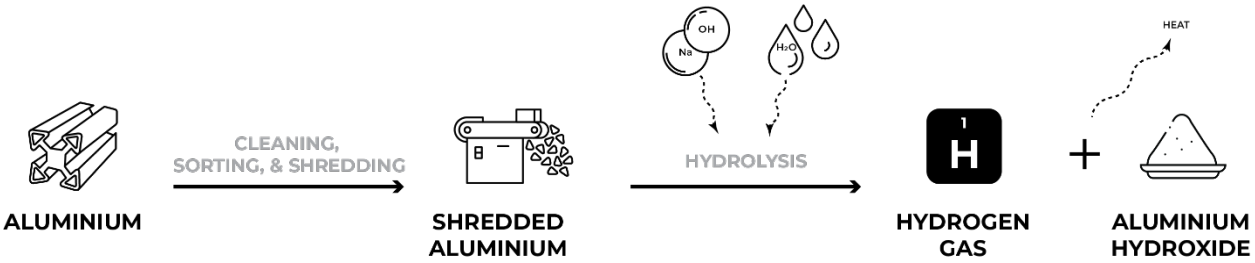


Figure 4-5: Illustrative Aluminium Hydrolysis Process (derived from Urbonavicius et al., 2024)

4.5.2 Proxy Foreground Processes

For foreground technologies that do not yet exist as standalone activities in ecoinvent, process data were implemented as "proxy" unit processes by cloning the closest available activity and minimally editing its exchanges to match the emerging route. Only the exchanges required to represent the new unit operation were edited (inputs, yields, and key utilities) while maintaining mass balances. Appendix B summarises the proxy construction for the two emerging technologies (MS-FSC and hydrolysis). Proxies for reuse, landfill, and conventional recycling routes follow the same minimum-change logic and are also documented in Appendix B.
Geographic Regionalization of Foreground Proxies

Utility regionalization was applied to proxy foreground processes for MS-FSC and aluminium hydrolysis by relinking electricity and other energy carriers to geography-consistent providers. A complete relinking log, including original providers and replacement markets, is provided in Appendix B.

4.5.3 Evidence-Based Scaling

Evidence-based scaling was used to convert the available MS-FSC and hydrolysis evidence into implementable C3–C4 foreground inventories for this case study. For each emerging route, a

single set of deterministic central values was instantiated under both the contemporary and prospective backgrounds to ensure comparability across scenarios. Parameters expected to vary with configuration, scale, or operational control were not embedded in the baseline and were instead evaluated explicitly in Step 5 as sensitivity/uncertainty levers (Thonemann et al., 2020; Tsoy et al., 2020).

4.5.3.1 Multi-Step Friction-Stir Consolidation (MS-FSC): Evidence-Based Scaling

MS-FSC was modelled as an emerging solid-state recycling route in which the dominant, defensible inventory driver is measured electricity demand per unit billet produced, while upstream preparation steps (shredding, degreasing/drying) are represented using transparent proxy-based unit operations held constant in structure across time horizons. Consistent with ex-ante LCA guidance, prospective scaling was implemented by preserving the same unit-process topology and mass/yield logic, then translating a small number of mechanistically interpretable operating parameters from a contemporary configuration to a prospective configuration, with the residual uncertainty treatment deferred to Step 5.

For MS-FSC, Ingarao et al. (2023) report phase-resolved power–time measurements that allow total electricity to be decomposed into two physically meaningful components: (A) productive consolidation energy and (B) transition (ramp/overhead) energy. The transition phase arises from test-equipment constraints during load ramping from pre-load to consolidation load, whereas consolidation energy corresponds to the productive work of bonding and densification. Because the experimental protocol is multi-step, the overall cycle comprises repeated transition–consolidation pairs; the reported phase energies were extracted on a 20 g billet basis and converted to specific intensities (kWh/kg) for use in the foreground unit process:

- ***Process A (Productive Consolidation):*** 0.267 MJ per 20 g (13.35 MJ/kg = 3.708 kWh/kg)
- ***Process B (Transition Overhead):*** 0.355 MJ per 20 g (17.75 MJ/kg = 4.931 kWh/kg)

Rather than applying a generic efficiency uplift to the full measured lab energy, the prospective central case treated the transition component (B) as a non-representative laboratory overhead. It is explicitly explained by the authors as equipment-limited load ramping, not a fundamental requirement of steady operation. In other words, the “improvement” is not an assumed learning curve applied to the physics of consolidation (A), but the removal (or strong reduction) of a measurement artifact whose cause is already diagnosed by domain experts in the source study. As a result, the central values for each baseline were selected as:

- ***Contemporary baseline (2025):*** Electricity demand is modelled as A + B (3.708 + 4.931 = 8.639 kWh/kg billet), representing the measured multi-step lab-scale total draw.

- ***Prospective Baselines (2050, SSP Backgrounds):*** Electricity demand is modelled as A only (3.708 kWh/kg billet) under an industrialized implementation framing, with (B) reserved for explicit sensitivity treatment.

These scale-up hypotheses are parameterized and tested explicitly in Step 5 as part of the uncertainty and sensitivity analysis.

4.5.3.2 Aluminium Hydrolysis: Evidence-Based Scaling

Aluminium scrap hydrolysis was modelled as a low-TRL EoL route whose prospective performance depends primarily on a small number of engineering-controlled levers rather than on incremental tuning of minor inputs. Scaling from the contemporary (2025) implementation to prospective (2050) implementations was therefore operationalized as a parameter translation that preserved the same unit-process structure, mass-balance logic, and service boundaries across years, while updating only the central values of a limited set of key parameters to reflect plausible maturation (Wender et al., 2014; Bergerson et al., 2020; Tsoy et al., 2020).

The hydrolysis route retained a consistent C3–C4 foreground definition in both time horizons: prepared aluminium scrap entered as the treated feed, alkaline liquor and water demands were accounted for, hydrogen conditioning burdens were represented via proxy services (e.g., PSA), and co-product quantities (H_2 and $Al(OH)_3$) were tracked for subsequent Stage D treatment rather than credited within the route itself. This ensured that differences between contemporary and prospective results were attributable to explicit scale-up assumptions rather than to changes in process topology.

Scaling was implemented through a compact set of variables that map directly to design intent and operating practice reported in experimental and early system studies (Urbonavicius et al., 2024). Central-value changes from 2025 to 2050 were applied as denoted in Table 4-5.

Table 4-5*: Prospective (2050) Central Scaling Parameters for Aluminium Hydrolysis (EoL Treatment)*

Scale Description	***Symbol***	***Central Value***	***Units***	***Interpretation***
Metallic Al fraction in prepared scrap	f_{Al}	1.00	kg Al / kg prepared scrap	Central-case assumes prepared scrap stream is effectively metallic; heterogeneity is handled via Step 5 ranges (not by lowering the central).
Effective Al conversion (reacted fraction)	X_{Al}	0.95	–	Represents improved process control (reduced passivation / better solids handling) without claiming best-case completion.

Working-solution intensity	L	150	L / kg Al (metal basis)	Engineered mixing and liquor management reduce dilution and total liquid handling versus lab-style operation.
NaOH molarity (working liquor)	C_{NaOH}	0.240	mol / L	Kept constant to avoid embedding "free" chemistry-performance gains; scale-up is represented mainly by liquor reuse and reduced dilution.
NaOH make-up fraction (recirculation proxy)	f_{makeup}	0.20	–	Represents partial liquor reuse (i.e., 20% make-up of the caustic inventory) without introducing a speculative NaOH regeneration unit.
Scrap preparation yield into hydrolysis feed	Y_{prep}	0.85	–	Modest improvement in preparation/handling losses and rejects relative to contemporary proxy practice.
PSA recovery (crude → usable H_2)	R_{PSA}	0.95	–	Reflects tighter operation/conditioning performance in an engineered PSA service assumption.
Auxiliary process electricity (mixing/pumping/solids handling)	E_{aux}	0.15	kWh / kg prepared scrap	Non-ideal but necessary electrical loads for an engineered reactor system (agitation, pumping, solids handling, controls).
Thermal management electricity (exothermic heat absorption/control)	E_{therm}	0.05	kWh / kg prepared scrap	Captures the energy penalty of managing the exothermic reaction heat (e.g., circulation through heat exchanger/cooling loop) so temperature remains in the operating window.

These scale-up translations defined the deterministic central configurations used to generate the prospective inventories for each scenario-consistent background.

4.5.4 Stage D Substitution

Stage D substitution was implemented for this case study as avoided production associated with net equivalent recovered products from the modelled C3–C4 routes, using the scenario-year marginal mixes defined in §4.4. For the aluminium curtainwall scope, Stage D accounting was applied after the route chain produced a creditable output (following the equivalence and conditioning narrative in §4.5.4.1) and was then scaled from the route basis to the curtainwall

functional unit (3.67 kg Al per m^2). Landfill generated no credit, whereas reuse, conventional recycling and both MS-FSC and aluminium hydrolysis were mapped to specific displaced products. The emerging technologies' mappings are detailed below and while reuse and conventional recycling can be found in Appendix D.

4.5.4.1 Product-to-Market Mapping

For the emerging routes, the displaced-product mappings were:

- ***MS-FSC:*** Qualifying output was mapped to extrusion-grade aluminium billet, or billet-equivalent feedstock, in the relevant Canadian receiving market. Avoided burdens were calculated using the scenario-year marginal mix for extrusion-grade billet defined in §4.4.

- ***Aluminium Hydrolysis:*** Two candidate outputs were assessed and credited independently once represented in their implemented product basis.

    1. Purified gaseous hydrogen, represented using the implemented low-pressure hydrogen market proxy, was mapped to the Canadian marginal hydrogen market in §4.4.

    2. Recovered aluminium hydroxide solids, $Al(OH)_3$, were mapped to the relevant background market used for aluminium hydroxide supply (global (GLO)/rest of world (RoW) as specified in the proxy implementation), with avoided burdens taken from the scenario-consistent background market because a custom marginal mix was not defined for this product class.

4.5.4.2 Equivalence & Conditioning

Functional equivalence was established only for routes that produced non-obvious outputs. Landfilling did not displace any product and therefore did not require equivalence. Reuse and conventional remelting were treated as established routes whose inventories already conformed to standard specifications for the displaced products, so no additional proof was required beyond the documented inventories. Formal equivalence was therefore solely demonstrated for MS-FSC and aluminium hydrolysis and modelled in the foreground up to the point where the recovered output met the displaced product's specification.

Multi-Step Friction-Stir Consolidation (MS-FSC)

Equivalence was demonstrated against extrusion-grade billet, or billet-equivalent feedstock, rather than against finished structural sections. This reflects the current solid-state recycling literature, which supports dense consolidated bulk outputs and billet-type intermediate products from aluminium chips, but does not justify a generalized claim of direct finished-section equivalence across applications. The equivalence screen therefore remained at the billet/feedstock level: sound consolidated form, defensible chemistry and impurity compatibility with the intended feedstock class, and routing to the Canadian receiving market defined for extrusion-grade billet.

Peer-reviewed MS-FSC studies indicate that dense, defect-lean billet-type outputs can be produced from aluminium chips and fine scrap when preparation and processing parameters are appropriately controlled, with microstructural continuity and mechanical indicators approaching wrought benchmarks (Wan et al., 2017; Latif, Ingarao, & Fratini, 2022; Latif, Ingarao, Gucciardi, et al., 2022; Ingarao et al., 2023; Demirci et al., 2025). On this basis, the deterministic baseline treated qualifying MS-FSC output as billet-equivalent for substitution purposes. In deterministic runs, all modelled qualifying output was credited at full realized yield; in probabilistic runs, reduced qualifying share was handled explicitly through pass-share uncertainty rather than being embedded implicitly in the deterministic baseline. The entire process and system boundary are depicted in Figure 4-6.

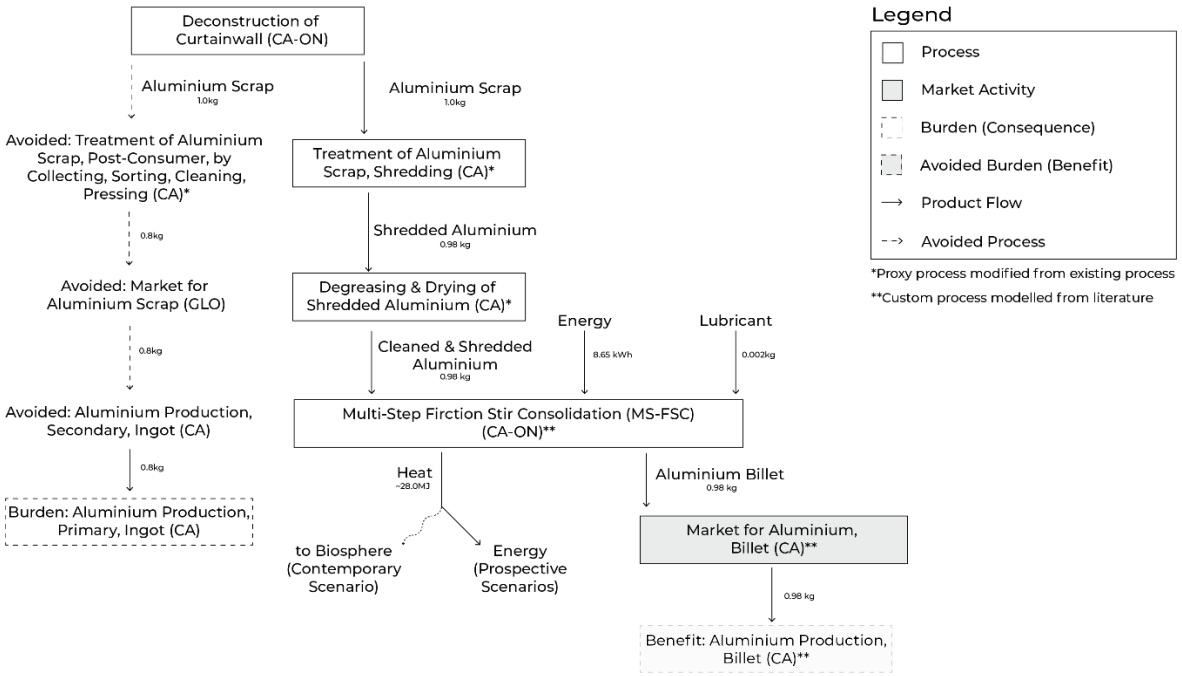


Figure 4-6*: Multi-Step Friction Stir Consolidation (MS-FSC) Contemporary System Boundaries & Crediting (Unit Process)*

Aluminium Hydrolysis

Hydrolysis generated two candidate products (see Figure 4-5), hydrogen and aluminium hydroxide, each assessed separately for Stage D substitution. The hydrogen stream was purified by pressure swing adsorption (PSA), and in the final implementation the credited receiving-system product was represented through a low-pressure hydrogen market proxy rather than a separate medium-pressure delivery basis. Under these conditions, the product was treated as equivalent to the gaseous hydrogen product represented by the low-pressure hydrogen receiving-system proxy in the Canadian marginal market defined in §4.4.

The aluminium hydroxide $Al(OH)_3$ stream was treated as a recovered solids output mapped to the selected background aluminium-hydroxide market proxy. In the final implementation, no

separate drying step was imposed as part of the credited product basis. Instead, reaction-associated water was assumed to evaporate or otherwise separate to the biosphere rather than being assigned as an added conditioning burden to the $Al(OH)_3$ product. The hydrolysis claim therefore remained conditional on the implemented product basis for each output rather than on any additional post hoc downscaling rule. The hydrolysis system boundary is illustrated in Figure 4-7.

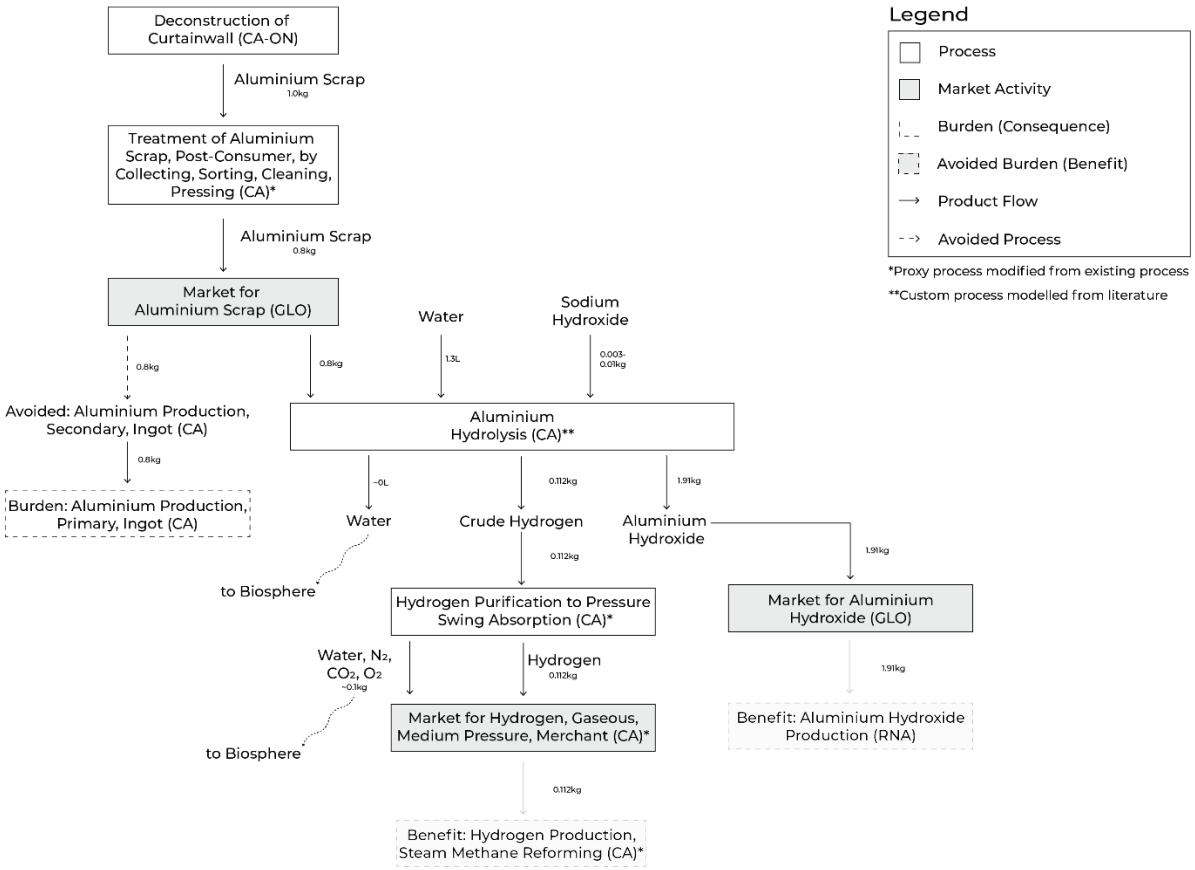


***Figure 4-7**: Aluminium Hydrolysis Contemporary System Boundaries & Crediting (Unit Process)*

For this aluminium case study, MS-FSC and hydrolysis required no additional Stage D downscaling or exclusions beyond the stated equivalence and conditioning assumptions.

4.5.5 Deterministic Results

This subsection synthesizes the deterministic EoL results into an interpretation that is defensible within the pCLCA structure implemented in this study. The central interpretive challenge in this case study is that net results for crediting routes are frequently governed by Stage D, while some routes, particularly those that displace services outside the metals sector (e.g., hydrogen), exhibit strong scenario dependence. These patterns are treated here as consequences of applying beyond-

the-system-boundary substitution logic to recovered outputs whose displaced counterfactuals reside in external markets, rather than as anomalies in the calculations.

In line with the Methods (§3.5), this subsection treats "deterministic" as a set of scenario-year point estimates produced using the Step 4 central foreground parameter set, with scenario differences arising from the frozen Step 2 backgrounds and Step 3 marginal mixes. Full interval reporting (p5–p95) and structural stress tests are reserved for Step 5.

Accordingly, this subsection distinguishes:

(i) what is directly supported by deterministic outputs (decomposition patterns, sign changes, scenario spreads in the reported terms),

(ii) what remains conditional on substitution governance defined earlier (counterfactual definitions, market scope, equivalence rules, marginal response representation), and

(iii) what must be demonstrated through targeted sensitivity and uncertainty analysis (Step 5).

Deterministic life cycle impact assessment (LCIA) results are reported for the aluminium sub-system represented in the study functional unit (FU). Climate change (ReCiPe 2016 Midpoint (H), climate change indicator) is used as the primary indicator for interpretation in this subsection (Figure 4-8; Table 4-6). Additional midpoint indicators were computed for burden-shifting screening and are summarized in Appendix E using the same decomposition logic.

All burdens and credits are temporally aligned to the scenario deconstruction year as defined in §4.2. No dynamic inventory timing, discounting, or delayed-credit factors were applied. Deterministic outcomes serve two roles: (i) establishing the reference points for Step 5 uncertainty propagation within each scenario and (ii) screening for high-leverage dependence channels that motivate targeted stress tests.

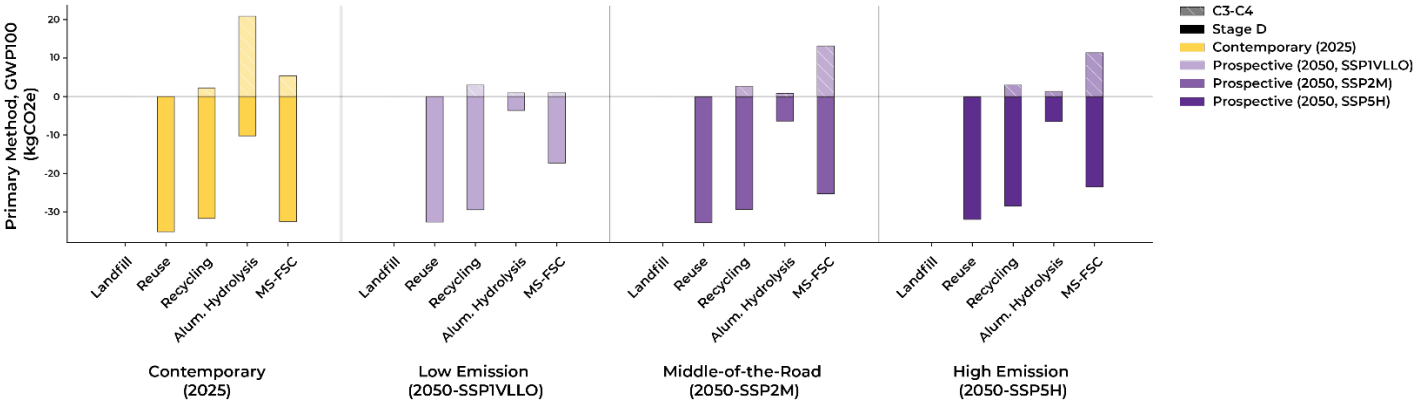


***Figure 4-8**: Decomposed deterministic climate change results (GWP100, kg CO_2-eq): C3–C4, Stage D, and net, across EoL routes for 2025 contemporary and 2050 prospective SSP variants*

***Table 4-6**: Decomposed deterministic climate change results (GWP100, kg CO_2-eq): C3–C4, Stage D, and net, across EoL routes for 2025 contemporary and 2050 prospective SSP variants*

Scenario	*Booking Year*	*EoL Route*	*C3–C4 (kg CO_2e)*	*Stage D (kg CO_2e)*	*Total (kg CO_2e)*
Contemporary	2025	Landfill	+0.095	0.000	+0.095
Contemporary	2025	Reuse	+0.011	-35.195	-35.184
Contemporary	2025	Conventional Recycling	+2.108	-31.729	-29.622
Contemporary	2025	Aluminium Hydrolysis	+20.886	-10.247	+10.638
Contemporary	2025	MS-FSC	+5.317	-32.521	-27.205
SSP1-VLLO	2050	Landfill	+0.076	0.000	+0.076
SSP1-VLLO	2050	Reuse	+0.007	-32.695	-32.688
SSP1-VLLO	2050	Conventional Recycling	+2.977	-29.464	-26.487
SSP1-VLLO	2050	Aluminium Hydrolysis	+0.946	-3.635	-2.689
SSP1-VLLO	2050	MS-FSC	+0.948	-17.294	-16.346
SSP2-M	2050	Landfill	+0.081	0.000	+0.081
SSP2-M	2050	Reuse	+0.008	-32.780	-32.772
SSP2-M	2050	Conventional Recycling	+2.702	-29.310	-26.608
SSP2-M	2050	Aluminium Hydrolysis	+0.822	-6.328	-5.507
SSP2-M	2050	MS-FSC	+13.156	-25.221	-12.064
SSP5-H	2050	Landfill	+0.077	0.000	+0.077
SSP5-H	2050	Reuse	+0.007	-31.914	-31.907
SSP5-H	2050	Conventional Recycling	+2.944	-28.432	-25.487
SSP5-H	2050	Aluminium Hydrolysis	+1.309	-6.449	-5.140
SSP5-H	2050	MS-FSC	+11.350	-23.460	-12.109

4.5.5.1 Stage D as a Qualified Substitution Claim

A central finding of the deterministic decomposition is that net EoL results for crediting routes are governed primarily by Stage D rather than by C3–C4 treatment burdens, meaning that net rankings for these routes are primarily rankings of substitution claims, not of treatment-chain burdens (Figure 4-8; Table 4-6). For the routes assessed here, C3–C4 burdens are consistently small relative to the magnitude of Stage D credits for the most credit-intensive pathways. For example, reuse shows C3–C4 burdens on the order of hundredths of a kg CO_2-eq (approximately 0.007–0.011 kg CO_2-eq across scenarios), while Stage D credits are on the order of tens of kg CO_2-eq (approximately −31.9 to −35.2 kg CO_2-eq). Conventional recycling exhibits the same general pattern, with C3–C4 in the low single digits (≈2.1–3.0 kg CO_2-eq) and Stage D again in the tens of kg CO_2-eq (≈−28.4 to −31.7 kg CO_2-eq). Even where treatment burdens are not negligible (notably, the contemporary hydrolysis case), the sign and magnitude of the net result remains tightly coupled to the Stage D term because the credit is a major component of the net balance. These dominance relationships are an expected feature of consequential modelling

whenever recovered outputs plausibly displace burdensome supply (e.g., primary aluminium production or fossil-derived hydrogen). However, it changes what is typically defined as "evidence" in the results and shifts the focus from how "dirty" an EoL technology is to how strong the avoided production (substitution) assumption is.

Within EN 15978-aligned module conventions, C3–C4 represents burdens from waste processing, conditioning, and disposal, whereas Stage D represents potential benefits and loads beyond the system boundary. In a consequential formulation, Stage D is not an emissions attribute of the technology. It is a counterfactual consequence that only materializes if a recovered output enters an external receiving system and displaces marginal supply within that system. The credibility of the Stage D term therefore depends on the credibility and transparency of the displacement definition. At minimum, a defensible substitution claim requires explicit specification of: (i) the displaced product or service, (ii) the receiving market scope, (iii) the marginal supplier response representation, and (iv) the substitution ratio or functional equivalence factor (Weidema et al., 1999; B. Weidema, 2003; Ekvall & Weidema, 2004). Put differently, Stage D is an emergent property of the EoL route plus the receiving system plus the claim governance rules that link them, not a property of the route in isolation.

This has three reporting and interpretation consequences for the remainder of this section. First, net values must be presented alongside decomposed values (C3–C4, Stage D, net) because decomposition is the primary device for preventing over-reading conditional counterfactual credits as if they were direct emissions outcomes. Second, the deterministic comparison should be read as a comparison of substitution claims whenever a dominance condition holds (for example, when Stage D substantially exceeds C3–C4), because small, plausible changes in claim definition can materially change net outcomes even if treatment burdens are stable. Third, the most decision-relevant robustness work in a Stage D-dominant setting is not limited to fine-tuning treatment parameters (transport distances, collection burdens, minor energy inputs). Instead, it prioritizes stress-testing the substitution structure itself, including displaced product definitions, market boundaries, marginal mixes, and substitution ratios.

Consistent with this claims-discipline stance, the deterministic baseline is accompanied by a compact substitution-claim ledger that makes the Stage D credit definition explicit and reviewable at the point of interpretation ensuring that each Stage D result can be traced to a specific displaced product/service, market scope, marginal response representation, and equivalence rule. Appendix A therefore summarizes the minimum claim elements underlying Stage D credits in the deterministic baseline—sufficient to interpret Stage D-dominant rankings as qualified displacement claims rather than emissions results. Where an element is held constant across scenarios (e.g., the displaced product definition), it is listed once at the route level, where it is scenario-dependent (e.g., marginal supplier representation for 2050 SSPs), it is recorded explicitly by year/scenario.

4.5.5.2 Interpreting Dominance

To interpret Stage D-dominant deterministic results defensibly, three dependence channels are distinguished because each changes net outcomes through a different mechanism:

1. ***Scenario Dependence*** *(between SSPs, holding modelling structure fixed):* Appears when Stage D (or net) shifts materially across SSP variants for the same route, indicating that the scenario-conditioned receiving system embedded in the counterfactual is influencing the credit magnitude, including receiving-system evolution (energy system transitions, industrial structure) and any scenario-driven changes to market representations. In this case study, scenario dependence is most clearly indicated where Stage D varies across SSPs relative to the route's own C3–C4 scale.
2. ***Structural Dependence*** *(within an SSP, changing the credit governance rules):* Structural dependence cannot be established from deterministic outputs alone, but dominance indicates high leverage: if Stage D dwarfs C3–C4, plausible alternatives for displaced-product proxies, market boundaries, equivalence ratios, or marginal response representation can move the net result substantially even when treatment burdens remain unchanged. This channel is directly supported when audits show that governance choices (e.g., localization of receiving-system proxies) measurably change the displaced system's unit burdens.
3. ***Parameter Dependence*** *(within a fixed structure):* Parameter dependence arises from numeric assumptions within the fixed governance structure (yields, energy intensities, conditioning requirements). The decomposition indicates where parameter refinement is decision-relevant: if C3–C4 is a residual term relative to Stage D, parameter refinement alone is unlikely to overturn conclusions unless the credit term is simultaneously stress-tested.

This framing supports an important restriction on deterministic interpretation: scenario-stable rankings are not sufficient evidence of robustness when Stage D dominates. Stability across SSPs can coexist with sensitivity to plausible structural alternatives within the same SSP. Consequently, this subsection treats dominance as a screen that prioritizes structural stress tests (Figure 4-9), it does not claim robustness on the basis of deterministic scenario sweeps.

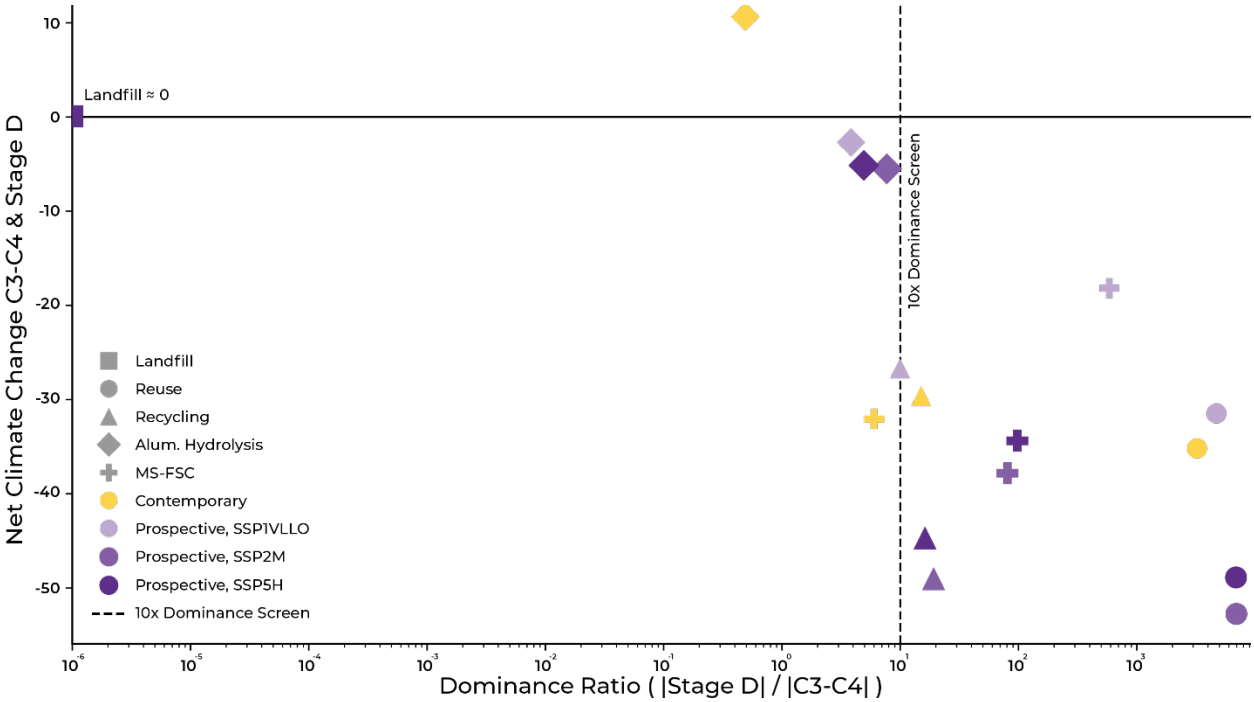


Figure 4-9: *Dominance Ratio vs Net Climate Change Interpretation Map*

Across routes, Figure 4-9 separates three groupings that guide the structure of the subsequent discussion. First, landfill anchors the non-credit baseline, confirming that net outcomes reduce to C3C4 burdens when no displacement claim is invoked. Second, within-sector metal displacement routes occupy a Stage D–dominant space, indicating that net results are largely governed by how the displaced aluminium production proxy is specified (market scope, supplier mix, equivalence), and by how that proxy evolves across scenarios. This motivates the next subsections' focus on credit governance audits and on scenario-conditioned receiving systems, rather than treating deterministic rankings as inherently robust. Third, emerging routes span a wider range of dominance, which flags that they cannot be interpreted uniformly: some configurations remain burden-governed while others become credit-governed once the displaced-system term becomes large relative to route burdens. This heterogeneity justifies treating emerging routes in a route-specific manner, emphasizing where structural dependence dominates interpretation and where parameter dependence remains decision-relevant.

Finally, the scenario dispersion in Figure 4-9 provides a bounded signal of scenario dependence without resolving structural dependence. Movement across SSPs indicates that the receiving system embedded in the Stage D counterfactual changes the magnitude of the credit, but dominance indicates the leverage of alternative claim structures within any single SSP.

4.5.5.3 Landfilling as a Baseline

Landfill consistently yields small positive deterministic results because it generates no substitution credit. While unsurprising, its interpretive value is non-trivial within this Stage D-

dominant results set. Landfill provides a baseline of pure treatment-chain burdens, against which the crediting routes can be interpreted.

This baseline serves two functions in the present results. First, it provides a sanity check that C3–C4 burdens behave as expected at the functional-unit scale (Figure 4-10): they remain positive and of plausible magnitude across the scenario variants. Second, it clarifies that net negativity is not an intrinsic feature of EoL processing, emerging only when an avoided-production claim is asserted. Maintaining this distinction reduces the risk of a common rhetorical error in circularity discussions where negative net results are treated as if they were "negative emissions," rather than conditional counterfactual consequences.

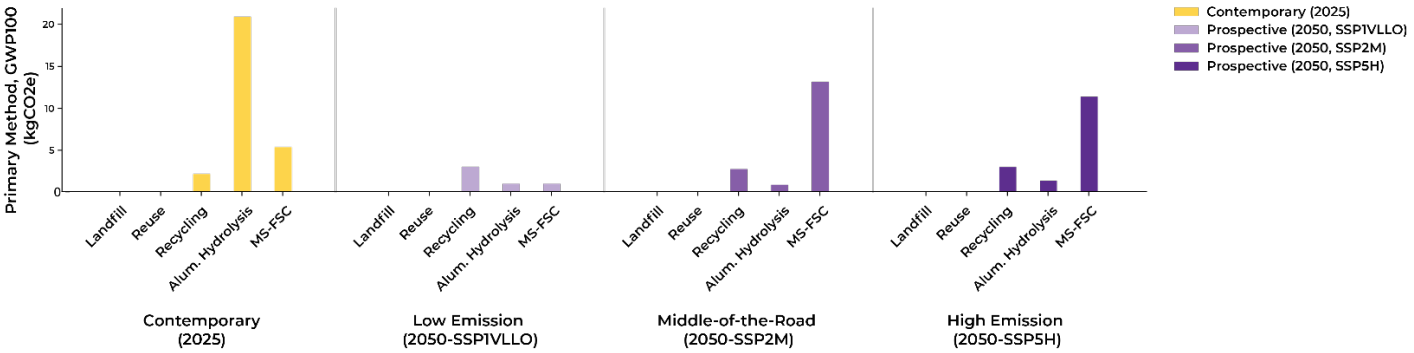


Figure 4-10*: Decomposed deterministic climate change results (kg CO_2-eq): C3–C4 only, across EoL routes for 2025 contemporary and 2050 prospective SSP variants*

4.5.5.4 Within-Sector Displacement: Reuse and Conventional Recycling

Reuse and conventional recycling exhibit the clearest Stage D dominance pattern in the deterministic baseline. Their net outcomes are largely governed by within-sector displacement of aluminium supply, while C3–C4 burdens remain comparatively small.

Two features are important for interpretation and for the Step 5 design:

- ***Scenario dependence is limited relative to credit magnitude*** *(screening outcome):* Across SSP variants, the Stage D term for these within-sector routes shows comparatively small variation relative to its magnitude. This supports a screening inference that scenario choice is not the primary driver of differences among these routes in the deterministic baseline.
- ***Structural dependence remains high leverage under dominance*** *(interpretive implication):* Limited scenario spread does not imply robustness because the comparison is still governed by a displacement claim. Under dominance, plausible within-scenario alternatives (displaced-product proxy aligned to recovered output specification, market boundary definitions, equivalence/substitution ratios, and marginal response representation) remain capable of materially shifting net outcomes.

For reuse and remelting-based recycling, the highest-leverage uncertainties are structural rather than parametric. Step 5 therefore prioritizes discrete displaced-product/market-scope/marginal-mix variants alongside (not instead of) parameter sampling of minor burdens.

4.5.5.5 Cross-Sector Displacement: Hydrolysis and Receiving-System Dependence

Among the assessed routes, hydrolysis is structurally distinctive because its primary credit is not an aluminium product but a cross-sector service (hydrogen supply), with a second credited co-product (aluminium hydroxide). This distinction is reflected in the decomposition and is central to the scenario-dependence story in the deterministic results.

Two deterministic signals make hydrolysis a priority hotspot for deeper investigation:

- ***Time/Representation-Driven Shift:*** For reuse, Stage D spans roughly −31.9 to −32.8 kg CO_2-eq across the 2050 SSPs (with C3–C4 remaining ~0.007–0.008 kg CO_2-eq). For recycling, Stage D spans roughly −28.4 to −29.5 kg CO_2-eq, again large relative to the C3–C4 burdens. This supports the screening inference that scenario choice is not the primary driver of differences among these within-sector routes in the deterministic baseline.
- ***Scenario-Conditioned Receiving-System Dependence:*** Compared with within-sector routes, hydrolysis exhibits a more visible SSP dependence in the Stage D term relative to the scale of its C3–C4 burden in 2050, indicating that the displaced hydrogen counterfactual meaningfully differs across prospective pathways. The deterministic implication is that scenario choice changes the displaced system, not merely the size of foreground burdens.

These signals also determine what evidence is most informative for hydrolysis in a deterministic LCIA. Because hydrolysis credits a cross-sector service, the Stage D magnitude is governed by the interaction between the credited output bundle delivered by the modelled hydrolysis chain per functional unit and the unit climate burdens of the displaced receiving systems used to value that bundle in each scenario-year.

Because the prospective hydrolysis foreground parameterization is held fixed within each scenario, the SSP dependence of Stage D in 2050 is interpreted primarily as receiving-system intensity dependence (the displaced system changes across SSPs), not as a change in credited output quantity.

To make this leverage explicit, Table 4-7 presents summaries of the two scenario-conditioned unit-intensity drawn directly from the receiving-system definitions used in Stage D mapping. First, it reports the displaced hydrogen intensity used in each scenario-year. Second, it reports the unit climate burden of grid electricity (medium voltage with high and low alternatives reported in for each scenario-year, which provides context for the pronounced time shift in the C3–C4 term for hydrolysis. Finally, rather than inferring an “implied credited hydrogen quantity” from Stage D, the table reports a break-even credit scaling factor. This diagnostic asks a simple question: by what multiplicative factor would the aggregate Stage D credit need to change for hydrolysis to

reach net zero, holding the foreground burdens fixed? Formally, the break-even scaling is computed as $s^* = \text{C3–C4}/|\text{ Stage D }|$. Values below one indicate that hydrolysis would remain net-beneficial even if the credit were made substantially more conservative whereas values above one indicate that the modelled credit would need to increase to offset burdens.

Table 4-7*: Hydrolysis receiving-system leverage and break-even credit scaling (derived from deterministic C3–C4, Stage D, displaced-hydrogen intensity, and unit electricity burdens used in each scenario-year).*

Scenario	***Year***	***Displaced H_2 intensity used (kg CO_2-eq/kg H_2)***	***Unit Electricity GWP (MV) (kg CO_2-eq/kWh)***	***Break-Even Credit Scalings****
Contemporary	2025	11.5895	0.5554	2.038
SSP1-VLLO	2050	1.9948	0.0263	0.260
SSP2-M	2050	11.7715	0.4161	0.130
SSP5-H	2050	12.6498	0.3606	0.203

The electricity diagnostic clarifies the time-driven collapse in the hydrolysis C3–C4 term. The unit medium-voltage (MV) electricity burden falls from 0.555 kg CO_2-eq/kWh in the contemporary case to 0.026 kg CO_2-eq/kWh in SSP1-VLLO, while remaining much higher in SSP2-M (0.416) and SSP5-H (0.361). This pattern provides direct deterministic evidence that background evolution is pathway-specific and large enough to plausibly compress electricity-linked portions of the hydrolysis foreground burdens in 2050, especially under SSP1-VLLO. In other words, before any credits are considered, the prospective backgrounds already create very different burden environments for an electricity- and energy-carrier-intensive treatment chain. Including the electricity unit intensities in Table 4-7 therefore serves a specific interpretive purpose: it shows that the observed reduction in C3–C4 is not an opaque modelling artifact but aligns with a large, scenario-conditioned change in a core upstream driver.

Within 2050, Table 4-7 then isolates the cross-sector receiving-system dependence that differentiates hydrolysis from within-sector routes. Because credited quantities are held constant across SSPs, the pathway contrast in Stage D reflects differences in the unit burdens used for the displaced hydrogen counterfactual (and any other displaced co-product systems, $Al(OH)_3$). The displaced hydrogen intensities show an especially strong contrast: SSP1-VLLO uses 1.995 kg CO_2-eq/kg H_2, while SSP2-M uses 11.772 and SSP5-H uses 12.650. This explains the qualitative Stage D pattern observed in the decomposition: SSP1 exhibits a smaller magnitude Stage D credit because the displaced hydrogen system is comparatively clean, so each unit of credited hydrogen displaces less GWP. Conversely, SSP2-M and SSP5-H assign much higher unit burdens to displaced hydrogen, which amplifies the credit term and yields larger net benefits, even though their electricity intensities (and therefore some portions of the hydrolysis C3–C4 burden) are higher than SSP1. Hydrolysis can therefore reach the same qualitative outcome across pathways through different combinations of burden suppression (via electricity and energy systems) and credit amplification (via receiving-system intensity).

The break-even scaling metric then bounds what can be claimed from deterministic results alone. In the contemporary case, the break-even factor exceeds two (2.038), indicating that hydrolysis would require roughly a doubling of the aggregate Stage D credit magnitude to break even and that under the modelled receiving-system definition, hydrolysis is robustly net burdensome in 2025. In contrast, all prospective cases have break-even factors well below one (0.130–0.260), meaning hydrolysis remains net beneficial even if the aggregate receiving-system credit were substantially reduced. The robustness ranking is also informative: SSP2-M is most robust (0.130), followed by SSP5-H (0.203), while SSP1-VLLO is comparatively more sensitive (0.260). This sensitivity is not a contradiction of SSP1's "clean" electricity and "clean" displaced hydrogen. It reflects the structural consequence of a cleaner displaced system: the per-unit credit is smaller, which leaves less buffer between Stage D and the reduced foreground burdens.

Taken together, Table 4-7 supports two deterministic conclusions. First, hydrolysis is the clearest receiving-system-sensitive route in the deterministic baseline because its net outcome is controlled by scenario-conditioned unit burdens in displaced cross-sector systems, rather than within-sector aluminium displacement alone. Second, once the unit electricity and displaced hydrogen intensities are surfaced explicitly, the next analytic step becomes subsequent robustness checks (Step 5) that target the above decision-relevant levers for hydrolysis. This includes testing conservative revisions to receiving-system intensities and by challenging energy-system assumptions that plausibly drive the collapse in C3–C4, then evaluating whether the net sign and route ordering persist.

4.5.5.6 Emerging Technology: MS-FSC and Compounded Dependence

MS-FSC is a stress case for the deterministic framework because it combines two features that amplify interpretive sensitivity. First, MS-FSC is represented as an emerging foreground process with performance that is not yet stabilized, so the deterministic result remains exposed to parameter dependence (e.g., yield, energy demand, and scale-up assumptions). Second, across both contemporary and prospective settings the net outcome is dominated by Stage D, which creates strong structural dependence. The sign and magnitude of the result become, in practice, a statement about the credited displacement definition embedded in the credit ledger, rather than a pure "process burden" comparison. When these two features coincide, deterministic interpretation requires explicit governance of what is being displaced and which receiving system is being invoked, because small structural choices in the receiving-system definition can dominate what otherwise looks like a technology comparison.

A first deterministic signal is Stage D dominance (credit dominance relative to treatment burdens). Table 4-8 shows that Stage D remains the larger-magnitude term and sets the direction of the total in all 2050 pathways, but the extent of cancellation differs sharply by scenario. In SSP1-VLLO, MS-FSC's C3–C4 burden is small relative to the credited displacement (0.948 vs −17.294 kg CO_2-eq), so the net remains primarily credit-governed (|Stage D|/|Total| ≈ 1.06). In SSP2-M and SSP5-H, however, C3–C4 becomes a first-order term (13.156 and 11.350 kg CO_2-

eq) that offsets roughly half of the credited magnitude (≈52% and 48% of |Stage D|, respectively). The net therefore becomes the residual of substantial cancellation (|Stage D|/|Total| ≈ 2.09 and 1.94), meaning that both the displacement claim (structural dependence) and MS-FSC performance parameters (parametric dependence) can plausibly shift conclusions. The contemporary case sits between these extremes: C3–C4 is non-trivial (5.317 kg CO_2-eq) but still materially smaller than the credited term (−24.736 kg CO_2-eq), leaving the total anchored by the displacement claim (|Stage D|/|Total| ≈ 1.27).

Table 4-8: *MS-FSC deterministic split and receiving-system context (primary method, GWP100).*

Scenario	***Year***	***Net (kg CO_2-eq/FU)***	***|Stage D| / |Net|***	***Unit MV electricity GWP (kg CO_2-eq/kWh)***
Contemporary	2025	−19.420	1.274	0.5554
SSP1-VLLO	2050	−16.346	1.058	0.0263
SSP2-M	2050	−12.064	2.091	0.4161
SSP5-H	2050	−12.109	1.937	0.3606

Within the 2050 pathways, the SSP spread in MS-FSC is not explained by Stage D alone. Stage D shifts from −17.29 kg CO_2-eq in SSP1-VLLO to −25.22 in SSP2-M and −23.46 in SSP5-H, but MS-FSC's own C3–C4 burdens also change substantially (0.95 in SSP1-VLLO versus 11.35–13.16 in SSP5-H/SSP2-M). As a result, the *net* outcomes partially converge despite different credits: SSP1-VLLO has the smallest magnitude credit but also the lowest treatment burdens (net = −16.35), whereas SSP2-M and SSP5-H have larger credits but much larger treatment burdens, yielding fewer negative nets (≈−12.06 to −12.11). This "offsetting" pattern is itself diagnostic as scenario-conditioned changes (most plausibly electricity-related differences that propagate into both the treatment chain and the displaced-aluminium proxy) can move both terms in the same direction, leaving the net as a residual that is highly sensitive to modelling structure and parameterization.

Electricity remains central to interpreting this behaviour because it operates in both parts of the balance. On the receiving-system side (Stage D), displaced aluminium production is electricity-coupled, so scenario-conditioned electricity burdens (and any voltage or market mapping choices used to represent them) provide a concrete mechanism for why the avoided-burden per unit displaced differs across SSPs. On the treatment-chain side (C3–C4), electricity is also a mechanistically defensible driver of MS-FSC burdens. Therefore, the prospective shift affects MS-FSC through two aligned channels: (i) a direct change in C3–C4 exposure through MS-FSC's electricity-linked burdens under each scenario's electricity representation, and (ii) a credit re-scaling in Stage D through scenario-conditioned burdens embedded in the displaced aluminium representation.

The fact that SSP2-M and SSP5-H exhibit very large |Stage D|/|Net| ratios (≈1.94–2.09) is informative. It indicates that the net is the difference between two large, electricity-coupled terms with opposite sign. In this regime, deterministic "route A beats route B" claims are fragile unless the displacement definition is explicitly governed and tested, because small changes to either (a) the displaced-aluminium proxy (market scope, product specification equivalence, voltage mapping) or (b) MS-FSC performance parameters (kWh/kg, yield, quality-gate pass rates) can shift the residual net materially.

MS-FSC therefore requires joint structural-plus-parametric tests. Parameter uncertainty on electricity demand and billet pass/fail (quality gate) is not meaningful unless paired with displaced-aluminium market/proxy variants (and the associated receiving-system electricity representation), because the net result remains credit-governed in sign and is cancellation-governed in magnitude in multiple 2050 SSPs.

4.5.5.7 Cross-Indicator Screening: Burden Shifting Under a Stage D-Governed Signal

Because climate-change results for several routes are governed by Stage D (i.e., the avoided-burden term dominated the net outcome), a minimum credibility check was performed to test whether the qualitative interpretation inferred from net GWP persisted when the same net ledger was evaluated across additional ReCiPe 2016 midpoint indicators. The screening was used to detect when "net-beneficial under climate change" acted as a misleading proxy for other decision-relevant mechanisms (e.g., toxicity, eutrophication, water use, resource scarcity), and to localize whether any divergence arose primarily from treatment-chain burdens (C3–C4) versus the imported burdens of the displaced receiving system (Stage D). Figure 4-11 summarizes this check as sign consistency versus GWP across the 17 non-climate midpoints for each of the 20 route-by-scenario cells. Cells near 17/17 indicated that GWP acted as a reliable qualitative proxy (net sign agreement) for most midpoint mechanisms under the same displacement ledger, whereas lower shares indicated conditions under which GWP was comparatively fragile as a qualitative signal.



Figure 4-11*:Heat map of Sign Consistency vs GWP for Non-Climate Midpoints*

Across the non-climate midpoint set, the climate-change signal was not qualitatively stable. Specifically, 9/17 midpoints exhibited at least one sign reversal relative to the GWP sign somewhere in the 20-cell grid, and 16/17 midpoints reordered routes (rank inversions) in at least one scenario. The largest sign instabilities occurred in ionising radiation (highest sign-change count), followed by freshwater and marine ecotoxicity, with additional sign reversals appearing in ozone depletion, eutrophication, water use, and mineral resource scarcity. This pattern was consistent with the central interpretive risk under Stage D dominance: a strong negative net GWP could coexist with burden shifting, so "net negative climate change" could not be interpreted as "overall better."

Screening is interpreted using three linked checks:

1. ***Alignment*** *(sign stability vs GWP):* Figure 4-11 showed that sign agreement with GWP was high for landfill, reuse, and recycling across scenarios (most cells ≥16/17), but systematically lower for hydrolysis in the 2050 scenarios (12/17). This indicated that the routes most reliant on cross-sector displacement were also those for which GWP least consistently represented other midpoint mechanisms.
2. ***Localization (C3–C4 vs Stage D):*** Where sign reversals occurred, they clustered into two qualitatively distinct modes:
    - ***C3–C4-Driven Reversals*** *(treatment burdens dominate):* Ecotoxicity indicators are the clearest case. For hydrolysis, freshwater and marine ecotoxicity reversals

in the 2050 scenarios are driven by large positive C3–C4 burdens (conditioning energy, chemicals, upstream supply chains) that overwhelm the Stage D credit observed under GWP. For MS-FSC, the contemporary reversal follows the same pattern, while the SSP cases additionally reflect receiving-system dependence because the Stage D term changes sign and contributes net burdens under ecotoxicity mechanisms.

- ***Stage D-Driven Reversals*** *(receiving-system dependence dominates):* Ionising radiation, water use, ozone depletion, eutrophication, and mineral resource scarcity show sign reversals that are predominantly Stage D-driven, meaning the directionality changes because the displaced production system is environmentally "different" in those mechanisms than it is for climate change (or because its importance is small enough that C3–C4 becomes decisive). In other words, the Stage D claim is not neutral. A displacement that yields a strong climate benefit may yield a weak, negligible, or even opposite signal for other impact pathways.

3. ***Scenario Amplification*** *(SSP sensitivity outside climate change):* Several non-climate indicators exhibited scenario-dependent shifts that materially altered net interpretation for specific routes, including cases that crossed zero. This confirmed that receiving-system evolution under SSPs could dominate qualitative outcomes for non-climate mechanisms even when the climate-change result appeared strongly net-beneficial.

Accordingly, the screening showed that Stage D-governed climate benefits did not imply a uniform environmental preference once other midpoint mechanisms were considered, and it reframed the decision-relevant question as where and why divergences arose (treatment chain versus receiving-system displacement). Complete indicator grids and the supporting sign- and rank-diagnostic outputs were reported in Appendix E, together with representative decomposed plots for the indicators exhibiting the largest qualitative instabilities.

4.5.5.8 Deterministic Stress-Test Suite Derived from Dominance Screening

The deterministic decomposition does not establish robustness, but it does specify which assumptions have the leverage to move net outcomes and therefore which tests are worth doing in Step 5. The dominance and scenario-spread patterns in Figure 1 and Figure 2 imply a small, targeted stress-test suite:

- ***Reuse and Conventional Recycling*** *(within-sector aluminium displacement):* Stress tests should prioritize substitution governance rather than minor treatment parameters. This includes: (i) displaced-aluminium proxy variants consistent with recovered output specification, (ii) alternative market boundary definitions (receiving scope), and (iii) alternative marginal response representations within the same scenario-year.

- ***Hydrolysis*** *(cross-sector displacement):* Tests should explicitly separate "SSP pathway effects" from "proxy/receiving-system definition effects." Practically, this means varying displaced-hydrogen counterfactual definitions that remain consistent with the scenario framing and challenging the receiving-system intensities that dominate Stage D (including electricity-market mapping variants where relevant).
- ***MS-FSC*** *(emerging + credit-governed):* Step 5 should use joint tests that vary (i) displaced-aluminium receiving-system definitions (the structural lever) and (ii) key foreground performance parameters (yield and electricity demand/scale-up assumption). The goal is not to produce a single "better" number, but to identify persistence/threshold conditions: what changes are required to flip sign or reorder routes.

4.5.5.9 What Deterministic Results Establish

The deterministic results establish three things that are directly supported by computed outputs:

1. ***Governance Dominance:*** For several routes, net outcomes are governed primarily by Stage D, so net rankings in those cases are primarily rankings of substitution claims rather than of treatment-chain burdens.
2. ***Receiving-System Sensitivity is Route-Dependent:*** Cross-sector displacement (hydrolysis) exhibits larger scenario-conditioned receiving-system dependence than within-sector displacement routes, consistent with its reliance on external service displacement.
3. ***Dominance is a Leverage Screen:*** Where Stage D dwarfs C3–C4, plausible alternative governance definitions (displaced proxy, market scope, marginal response, equivalence) have high interpretive leverage even if treatment burdens are stable.

The deterministic results do **not** establish robustness. Robustness claims require Step-6 testing of the structural and parameter levers implied by dominance and by scenario spreads.

4.5.5.10 Implications for Prospective Module D Reporting and Interpretation

A broader implication of these results is that standard Module D reporting conventions can encourage overconfident interpretation by presenting net benefits without explicit claim governance—especially in prospective studies where receiving systems are modelled to evolve.

In a pCLCA context, a Stage D credit should always be interpreted as conditional on:

- a specified receiving-system state (scenario and year),
- a specified displaced product/service definition,
- a specified marginal response representation, and
- an explicit substitution ratio / equivalence rule.

The reporting architecture used here, scenario-aligned receiving systems coupled with explicit substitution logic, operationalizes this discipline by requiring that each Stage D term be traceable to its displaced definition and receiving system. Decomposition reporting (C3–C4 / Stage D / net) and dominance metrics act as a governance mechanism that prevents conditional counterfactuals from being misread as direct emissions outcomes.

4.6 Step 5: Uncertainty Propagation and Sensitivity Evidence

This subsection provides the minimum robustness evidence needed to interpret the deterministic LCIA results as qualified consequential claims, without reproducing the full uncertainty and sensitivity analysis (reported in full in Appendix F). The previous step established that many net outcomes are Stage D governed and that some routes (notably hydrolysis) exhibit strong receiving-system sensitivity across SSPs. The following uncertainty propagation tests whether those qualitative interpretations persist once within-case variability is introduced, focusing on two decision-relevant questions: (i) whether the net outcome is sign-robust under the p5–p95 interval, and (ii) whether dominance regimes implied by Step 4 persist in distribution. In addition, Step 5 includes a separate sensitivity-screening exercise used to identify which sampled foreground levers are most strongly associated with net-outcome variation within fixed pathway backgrounds. The two workflows are related but distinct: uncertainty propagation characterizes score distributions under the implemented uncertainty design, whereas sensitivity screening ranks influential levers within that design.

4.6.1 Executed Uncertainty Layers and Interpretation Scope

Uncertainty propagation was executed under a deliberately limited design aligned to what was retained in the final implementation. For the contemporary (2025) case, background exchange uncertainty only was propagated via Monte Carlo sampling of embedded exchange uncertainty distributions while the foreground route structure, parameter values, and substitution-claim governance were held fixed. For the prospective (2050) case, foreground parameter uncertainty only was propagated for the emerging routes within each fixed SSP background. Prospective background exchange uncertainty was not propagated for the IAM-coupled databases, and no prospective joint background-plus-foreground Monte Carlo was retained. Accordingly, 2050 uncertainty should be interpreted as the combination of within-pathway foreground variability and between-pathway scenario spread, not as full within-scenario background-plus-foreground variance propagation. The deterministic results in Step 4 remain the fixed reference cases for interpretation, but the p50 values reported here are medians of separate stochastic runs and are therefore not expected to numerically reproduce the deterministic point estimates exactly. At the input stage, the foreground uncertainty runs use independent marginal sampling over the retained levers; however, those sampled draws are propagated through the shared foreground model and repeated Brightway LCI/LCIA evaluation, so the resulting outputs are not independent but structurally dependent model outcomes.

4.6.2 Robustness Screen: Sign Stability and Dominance Structure

Table 4-9 consolidates the primary robustness screen for GWP100 while Figure 4-12 and Figure 4-13 depict the net distributions across the contemporary and prospective routes, respectively. In the contemporary case, the within-sector displacement routes (reuse and conventional recycling) remain sign-robust, net-beneficial under p5–p95 and remain firmly Stage D governed, confirming that their net rankings are primarily rankings of qualified substitution claims rather than comparisons of treatment burdens. Landfill remains a stable positive baseline (no Stage D). MS-FSC remains sign-robust net-beneficial but occupies a mixed dominance regime, where Stage D governs sign while C3–C4 remains material for magnitude, indicating that both burden-side and credit-side assumptions remain decision-relevant for interpreting MS-FSC beyond a purely "credit-driven" narrative. Hydrolysis is the contemporary robustness hotspot. Its net distribution straddles zero under p5–p95, so the contemporary configuration cannot support a categorical claim that hydrolysis is either net-beneficial or net-burdensome when background exchange variability is considered.

In 2050, foreground-only uncertainty shows that pathway choice governs whether hydrolysis admits a single-sign narrative under uncertainty. Hydrolysis remains sign-uncertain under SSP1-VLLO and SSP2-M (p5–p95 crosses zero), but becomes sign-robust, net-beneficial under SSP5-H. Importantly, this pathway sensitivity is expressed primarily as a shift in the net balance relative to zero rather than a reversal of component directions. Across SSPs, median C3–C4 remains positive while median Stage D remains negative, but changes in their relative magnitudes move the net distribution across the sign boundary. In contrast, MS-FSC remains sign-robust net-beneficial across all SSPs under the modelled scale-up uncertainty levers, with dominance ratios indicating Stage D continues to anchor sign and that pathway sensitivity manifests mainly as magnitude shifts. These uncertainty results should therefore be read as conditional on the implemented uncertainty design: contemporary variability reflects background exchange uncertainty around a fixed foreground implementation, whereas prospective variability reflects foreground parameter uncertainty around a fixed SSP-conditioned background.

***Table 4-9**: Robustness Screen for Uncertainty Propagation (GWP100; medians and p5–p95; dominance at p50)*

Route	***n***	***Net p50***	***Net p5***	***Net p95***	***Sign robust? (p5–p95)***	***|D|/|C3–C4| (p50)***	***Dominance Structure***
Panel A — Contemporary (2025): Background Exchange Uncertainty							
Landfill	1600	0.109	0.082	0.149	**Yes**	–	No Stage D
Reuse	1600	−51.140	−60.100	−44.500	**Yes**	4.08e+03	Stage D dominated
Conventional recycling	1600	−45.386	−55.222	−37.642	**Yes**	21	Stage D dominated
Hydrolysis	2200	12.699	−0.150	32.103	**No**	0.46	Burden-dominant
MS-FSC	1400	−26.035	−32.484	−20.589	**Yes**	4.8	Mixed dominance
Panel B — Prospective (2050): Foreground-only Uncertainty (conditional on fixed SSP background)							
SSP1-VLLO: Hydrolysis	5200	+0.579	−1.362	+4.041	**No**	0.86	Near-cancellation
SSP2-M: Hydrolysis	2000	−2.412	−4.239	+0.733	**No**	1.67	Credit-leaning (sign-uncertain)
SSP5-H: Hydrolysis	2600	−3.044	−4.315	−0.821	**Yes**	2.02	Credit-dominant (sign-robust)
SSP1-VLLO: MS-FSC	1400	−19.274	−21.294	−17.250	**Yes**	36.10	Stage D dominated (extreme)
SSP2-M: MS-FSC	1400	−31.933	−34.859	−29.017	**Yes**	5.67	Stage D dominated (material burden share)
SSP5-H: MS-FSC	1400	−29.312	−31.664	−26.763	**Yes**	5.94	Stage D dominated (material burden share)

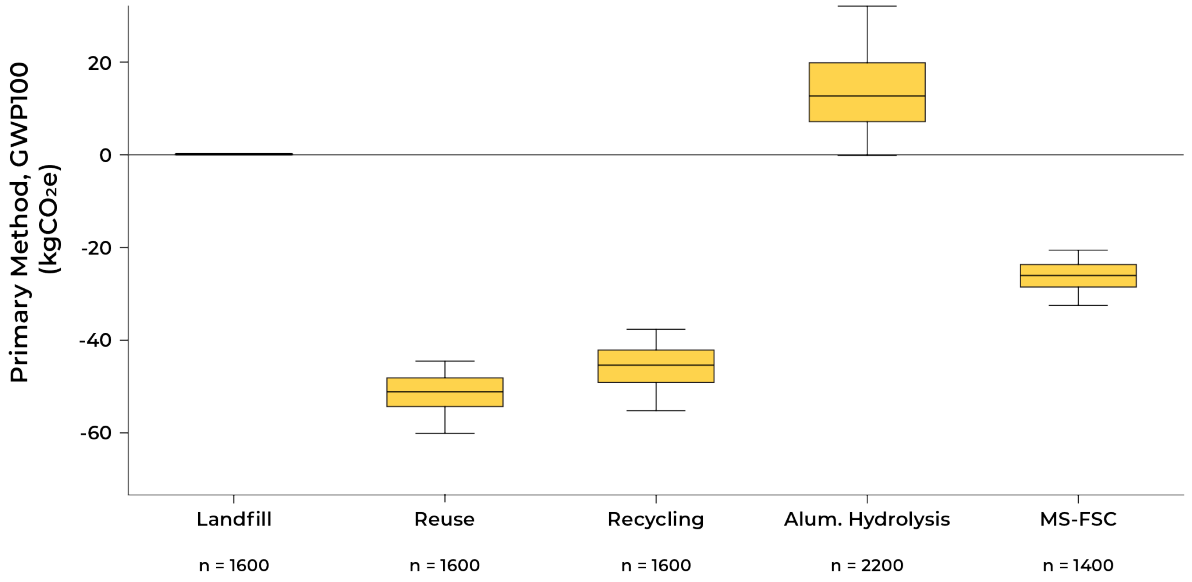


Figure 4-12: Contemporary background exchange uncertainty: net robustness across routes (GWP100)

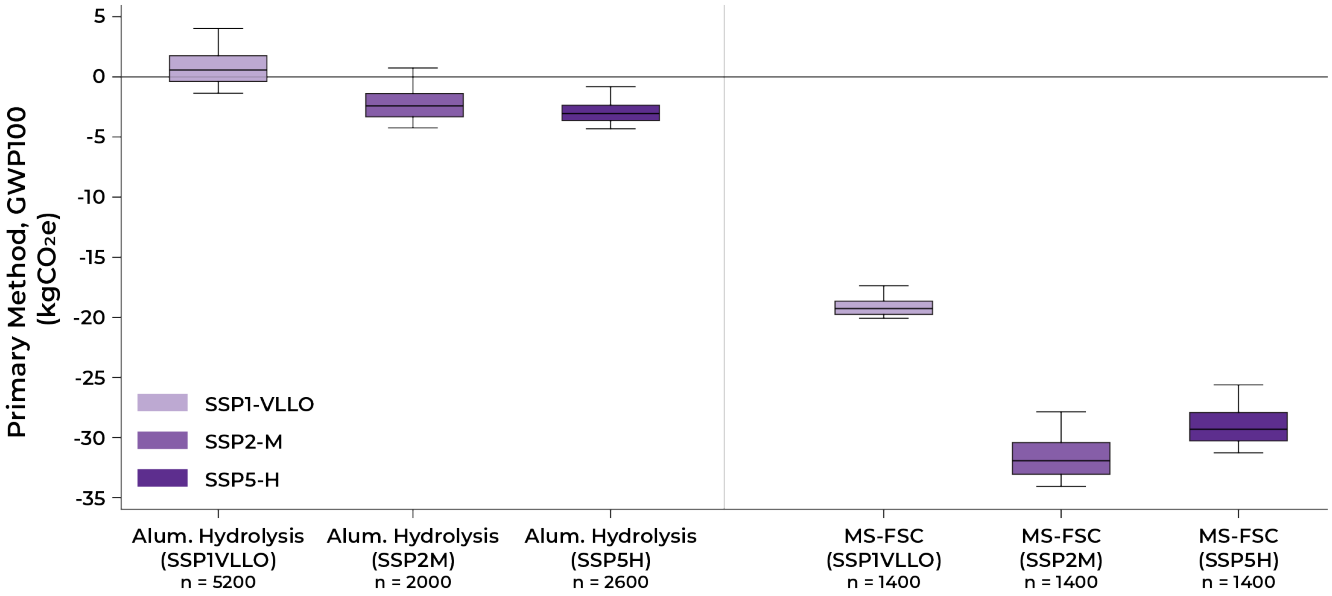


Figure 4-13: Prospective foreground exchange uncertainty: net robustness across routes and scenarios (GWP100)

4.6.3 Convergence and Adequacy

Because Monte Carlo uncertainty bands can be sensitive to sample size, especially when the reported net outcome is the residual of two opposing terms (Net = C3–C4 + Stage D), convergence was evaluated using a distribution-stability diagnostic rather than a fixed iteration count.

At periodic checkpoints m (e.g., every fixed number of additional draws), the accumulated Monte Carlo samples for the net outcome were summarized using a fixed grid of empirical quantiles:

Equation 4-8

$$q^{(m)} = \left(\hat{Q}^{(m)}(0.05), \hat{Q}^{(m)}(0.10), \ldots, \hat{Q}^{(m)}(0.95)\right),$$

where $\hat{Q}^{(m)}(p)$denotes the empirical p-quantile estimated from all samples collected up to checkpoint m. This quantile vector provides a compact "fingerprint" of the p5–p95 distribution (its location, spread, and tail behaviour within the reporting band). In practical terms, the stopping rule tests whether the reported distribution has stabilized, rather than whether an arbitrary number of Monte Carlo draws has been completed.

To quantify how much the simulated distribution changed between two successive checkpoints m and m', the root-mean-square change was computed across the quantile grid:

Equation 4-9

$$qRMSE(m, m') = \sqrt{\frac{1}{K}\sum_{p\in\mathcal{P}}\left(\hat{Q}^{(m')}(p) - \hat{Q}^{(m)}(p)\right)^2}, \mathcal{P} = \{0.05, 0.10, \ldots, 0.95\},$$

where $K = |\ \mathcal{P}\ |$. This yields an absolute stability measure in the same units as the results (kg CO_2-eq per FU). It answers, on average, how much the p5–p95 curve moved between checkpoints.

Since absolute shifts are not directly comparable across routes with very different net magnitudes, we normalized the qRMSE by the absolute median net at the later checkpoint:

Equation 4-10

$$qRMSE_{rel}(m, m') = \frac{qRMSE(m,m')}{|\hat{Q}^{(m')}(0.50)|}.$$

This relative form makes the stopping rule interpretable across cases as it measures distribution drift as a fraction of a typical net value.

Convergence (early stopping) was declared when $\mathrm{qRMSE}_{\mathrm{rel}} < 0.01$(i.e., <1%) for multiple (3) consecutive checkpoints ("stable hits"), ensuring that the p5–p95 band had stabilized rather than briefly dipping below the threshold by chance. The achieved iteration counts are therefore part of the evidence base. In particular, near-zero regimes impose stricter requirements under this normalization. When $|\ \hat{Q}(0.50)\ |$ is small (near-cancellation), even modest absolute movement in the estimated quantiles can produce a large $\mathrm{qRMSE}_{\mathrm{rel}}$. This is why SSP1-VLLO hydrolysis required substantially larger budgets (e.g., $n \approx 5200$) to meet the same tolerance, whereas MS-FSC converged rapidly across SSPs (e.g., $n \approx 1400$) consistent with strongly signed net

distributions and a low-dimensional foreground uncertainty design. The convergence criterion is therefore part of the robustness evidence itself. It shows that the reported percentile bands are numerically stable enough to support sign-robustness statements.

Rather than fixing an arbitrary budget (e.g., $n = 1000$), the use of $qRSME_{\text{rel}}$ as a stability rule ensure that the *reported percentile band* (p5–p95) have actually stabilized for each route–scenario case. This matters because some nets are near-cancellation, where small denominators make percentile estimates drift longer and can change whether p5–p95 crosses zero. Using a convergence-based stopping rule therefore makes sign-robustness statements auditable and prevents comparing routes with unequal numerical stability.

***Table 4-10**: Convergence evidence (full diagnostics in Appendix F):*

Route	***SSP***	***n at stop***	***qRMSE_rel***	***Interpretation***
Hydrolysis	SSP1-VLLO	5200	0.00909	Near-zero net → stringent relative stabilization
Hydrolysis	SSP2-M	2000	0.00422	Further from zero → faster stabilization
Hydrolysis	SSP5-H	2600	0.00169	Strongly signed → stable quantiles
MS-FSC	All SSPs	1400	~0.00085–0.00138	Strongly signed → rapid stabilization

4.6.4 Sensitivity Hotspots

Sensitivity analysis was used to identify which uncertain foreground assumptions exert the greatest leverage on net outcomes within fixed SSP backgrounds, and to connect uncertainty patterns back to the deterministic dominance screening. This sensitivity work should be interpreted as a global foreground screening exercise rather than as a one-factor-at-a-time perturbation or a formal variance decomposition. In each iteration, the selected levers were sampled simultaneously from their marginal distributions, patched into the route model, propagated through a full LCI/LCIA recalculation, and recorded as one input–output realization. Influence was then summarized using three screening metrics: Spearman's ρ for monotonic association, Pearson's r for linear association, and a decile-effect contrast comparing mean outputs in the top versus bottom deciles of each lever. The ranking is therefore association-based and screening-oriented; it should not be interpreted as a Sobol-type attribution of variance.

Hydrolysis exhibits multi-lever control, with the strongest monotonic responses associated with hydrogen purification recovery (*R_PSA*) and caustic make-up fraction (*f_makeup*). These levers directly scale the credit-side and burden-side terms, respectively, making them sign-relevant whenever the pathway-conditioned net approaches zero (as in SSP1-VLLO, and to a lesser extent

SSP2-M). Because hydrolysis maps sampled primitive inputs into derived process quantities before exchange updates, the screening retains the internal structure of the encoded process model rather than simply perturbing unrelated exchanges independently. MS-FSC is dominated by credit realization / eligibility (*pass_share*), which scales the Stage D term nearly linearly. Transition overhead retention (*f_transition*) provides a secondary burden-side modulation that is typically magnitude-relevant but not sign-determining under the tested conditions. In contrast to hydrolysis, MS-FSC sensitivity is propagated through a simpler structural model centred on transition burden and Stage D credit passage; thus, its driver rankings should be read as route-logic sensitivities rather than as a detailed process-physics decomposition. These sensitivity patterns motivate persistence-condition reporting: explicitly stating the parameter and governance conditions under which qualitative conclusions persist or fail, rather than relying solely on deterministic route rankings in Stage D governed regimes. The resulting driver relationships across SSP backgrounds are illustrated in Figure 4-14 for hydrolysis and Figure 4-15 for MS-FSC.

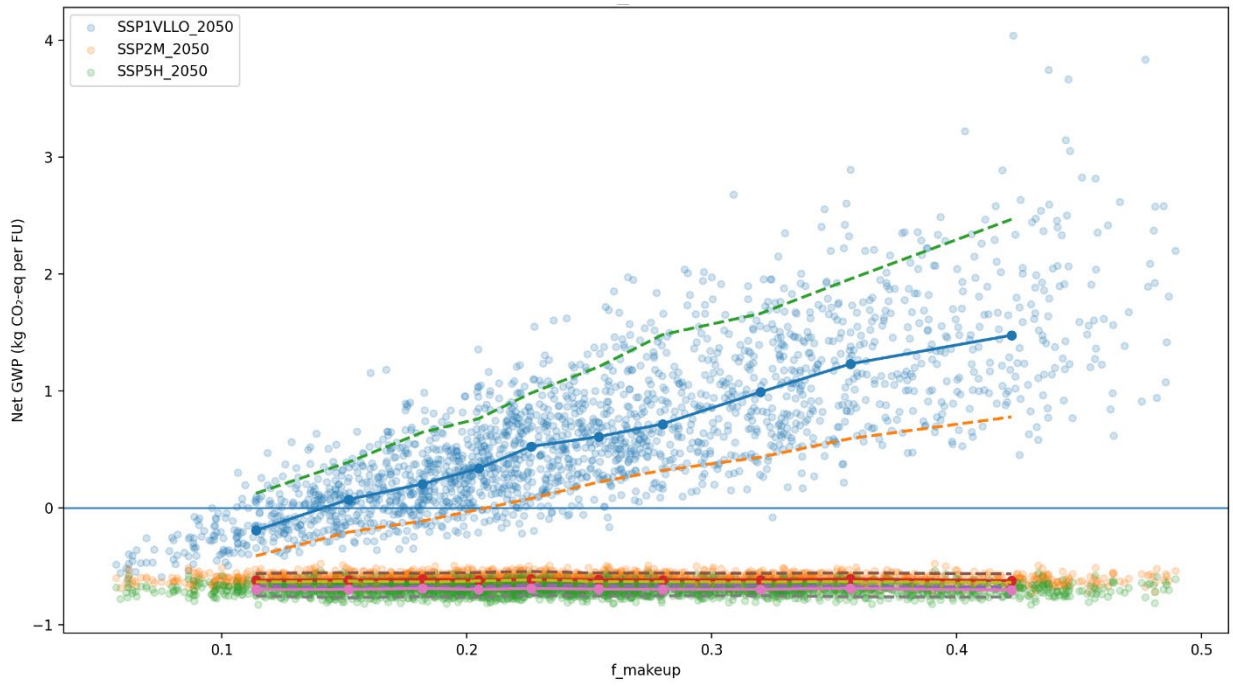


***Figure 4-14**: Hydrolysis driver plot showing the relationship between NaOH recirculated makeup fraction (f_makeup) and net outcomes across SSP backgrounds (Monte Carlo draws shown with SSP overlays).*

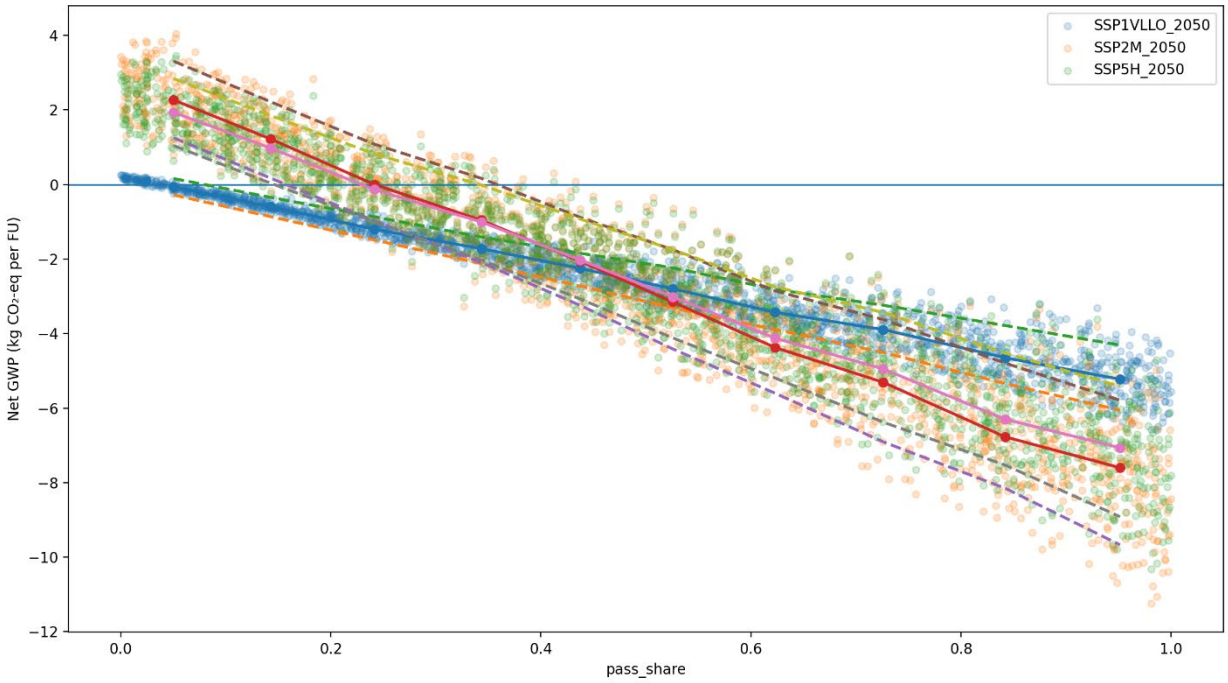


***Figure 4-15**: MS-FSC driver plot showing the relationship between credit realization share (pass_share) and net outcomes across SSP backgrounds.*

4.6.5 Implications for Interpreting Deterministic Results

The uncertainty propagation and sensitivity analysis support three interpretations in relation to the deterministic results. First, dominance is structural and persists under uncertainty. For within-sector displacement routes, Stage D governance remains the primary determinant of net outcomes even once uncertainty is considered, reinforcing that deterministic rankings in credit-dominant regimes are rankings of qualified substitution claims. Second, hydrolysis is the robustness hotspot across timeframes since it signals sign-instability in 2025 under background exchange uncertainty and remains sign-unstable in near-cancellation prospective configurations (SSP1-VLLO, SSP2-M), requiring probability-style reporting and targeted stress tests of both performance parameters and receiving-system definitions. Third, convergence diagnostics are integral evidence in cancellation regimes: iteration budgets are not transferable across routes or scenarios under relative stability diagnostics, and near-zero configurations require explicitly justified Monte Carlo stopping criteria to support tail reporting and sign claims. More broadly, the combined Step 5 evidence should be read as a conditional robustness layer: uncertainty propagation shows how far route conclusions persist under the retained uncertainty design, while sensitivity screening shows which foreground assumptions most strongly condition that persistence.

Full uncertainty propagation results, parameter range definitions, edge-case safeguards, and complete sensitivity grids are provided in Appendix F.

Chapter 5

5 Conclusions

This thesis developed and demonstrated a prospective consequential LCA (pCLCA) framework for evaluating end-of-life (EoL) strategies in the built environment under evolving background systems and explicitly governed substitution claims. Its principal contribution is the framework itself, with the case study serving as a demonstrative application used to operationalize, test, and interpret the framework under a realistic built-environment decision context. The central motivation, established in Chapter 1, is that built-environment decisions are increasingly justified using "net" results that combine treatment burdens (modules C3–C4) with beyond-the-system-boundary credits (Stage D), even though Stage D credits represent conditional consequences of market displacement rather than intrinsic emissions attributes of EoL technologies. Addressing this interpretive risk requires a workflow that is temporally consistent, transparent about displacement assumptions, and able to test whether conclusions persist under plausible future conditions and uncertainty.

To meet this need, the proposed framework couples scenario-consistent backgrounds derived from IAM pathways with a Canada-anchored marginal-supplier protocol, auditable foreground unit-process modelling for conventional and emerging EoL routes, and Stage D substitution rules that enforce like-for-like equivalence (including conditioning burdens and losses) prior to crediting. The case study applied this method to a Canadian curtain wall functional unit, tracking aluminium EoL pathways in a contemporary (2025) baseline and prospective (2050) scenarios under three IMAGE-derived pathways (SSP1–VLLO, SSP2–M, SSP5–H), thereby addressing a gap highlighted in the introduction: the lack of pCLCA studies that treat EoL dynamics and secondary markets as first-order drivers.

5.1 Methodological Contributions

The analysis advances pCLCA practice in three connected ways that align with the objectives stated in Chapter 1. First, it provides an end-to-end procedure for integrating IAM-driven background change into consequential modelling at a defined assessment year. Scenario-consistent backgrounds are constructed and frozen prior to downstream modelling, ensuring that foreground assumptions and substitution claims are evaluated against internally coherent future system states rather than static "today" inventories. Second, it formalizes a transparent marginal-supplier and market-boundary protocol tailored to a Canadian decision context. Rather than relying exclusively on tool-generated consequential markets, the approach defines displaced product markets through explicit equivalence and geography rules and computes marginal supplier mixes as a foreground construct where Canada-specific resolution is decision-relevant. This directly supports the thesis objective of assessing how supplier competitiveness and marginal production shifts influence results, while keeping overrides bounded, auditable, and reproducible. Third, the framework embeds uncertainty as an evidence layer rather than an

afterthought. Scenario uncertainty is represented through contrasts across IAM-derived pathways, while parameter uncertainty is propagated where it is defensible and interpretable. Convergence checks and sensitivity analysis are treated as requirements for reporting in configurations where net results are residuals of offsetting terms. Together, these contributions position pCLCA net results as governed consequential claims that must be interpreted against explicit receiving-system definitions, substitution rules, and uncertainty evidence.

5.2 Key Findings

The deterministic results establish a core interpretive pattern that directly addresses the "net-result" challenge raised in Chapter 1. For several crediting routes, net outcomes are governed primarily by Stage D rather than by C3–C4 treatment burdens. In these credit-dominant regimes, route rankings are primarily rankings of qualified displacement claims defined by displaced-product choices, market boundaries, marginal supplier representations, and equivalence rules. This is an expected property of consequential modelling when recovered outputs plausibly displace burdensome supply. However, it shifts what constitutes defensible evidence and negative net values should be interpreted as conditional displacement consequences, not "negative emissions," and must be presented alongside decomposition (C3–C4, Stage D, net) and claim documentation.

The case study also demonstrates that receiving-system sensitivity is route dependent, reinforcing the need for prospective modelling in EoL assessment. Hydrolysis is structurally distinct because its dominant credited service is cross-sector hydrogen rather than in-sector aluminium. Consequently, its Stage D term inherits scenario-conditioned variation in the displaced hydrogen system embedded in each IAM pathway's background and market representations. MS-FSC occupies a different, complementary position as an emerging EoL route with credit-governed net outcomes. It combines structural dependence (what is displaced and where) with parameter dependence (scale-up and eligibility assumptions), particularly in scenarios where net results reflect cancellation between electricity-coupled burdens and avoided production.

Uncertainty and sensitivity evidence strengthens and qualifies these deterministic interpretations in the precise way envisioned in the introduction. In the contemporary baseline, Monte Carlo propagation of background exchange uncertainty confirms that dominance patterns largely persist in distribution. Reuse and conventional recycling remain sign-robust and credit-governed, while landfill remains a stable positive baseline without Stage D. Hydrolysis emerges as a robustness hotspot in 2025 because its net can approach or cross the sign boundary under background exchange variability, indicating that categorical claims of benefit or burden require probability-based reporting and targeted stress tests. In the prospective setting, uncertainty is conditional on fixed SSP backgrounds because exchange-level background uncertainty is not consistently preserved or defensibly transformable after IAM-coupling transformations. Within this conditional framing, hydrolysis remains pathway-sensitive, shifting between sign-uncertain

and sign-robust configurations depending on the receiving-system environment. MS-FSC remains sign-robust under the tested scale-up levers, but its magnitude remains governance-conditional because realized Stage D credit depends on eligibility and substitution realization.

Finally, cross-indicator screening supports a broader implication for practice: “net-beneficial under GWP” does not imply uniform environmental preference across mechanisms. Under Stage D-dominant climate signals, burden shifting and receiving-system differences can drive sign changes or rank inversions in non-climate midpoint categories. This reinforces the thesis’s claims-discipline stance that Stage D is not a neutral accounting add-on and that robust interpretation requires explicit governance, decomposition-based reporting, and cautious generalization beyond the primary climate indicator.

5.3 Scope and Limitations

Several limitations follow directly from the scope and from current constraints of IAM-to-LCI pipelines, consistent with the limitations outlined in Chapter 1. The prospective backgrounds were generated as single-year 2050 snapshots via premise and frozen for reproducibility, but only electricity and fuel/energy-carrier systems were structurally transformed. Other sectors retained contemporary structural configurations and therefore represent a structural limitation in long-run representations. Spatial resolution was inherited from IAM coupling (e.g., regional markets) rather than fully Canada- or province-specific supply chains, and Canada-specific consequential interpretation was implemented primarily through foreground market definitions and marginal supplier construction rather than through comprehensive background disaggregation. In addition, prospective exchange-level background uncertainty was not propagated because uncertainty metadata are not consistently retained or defensibly transformed after scenario rescaling; prospective uncertainty bands therefore represent foreground variability around fixed SSP inventories rather than full within-scenario LCI uncertainty. Moreover, the foreground uncertainty design samples uncertain primitives primarily from independent marginal distributions unless dependence is imposed through explicit route equations. The resulting approach is methodologically sound as a Monte Carlo propagation and global screening framework, but it should not be interpreted as a fully correlated physical uncertainty model or as a full first-order variance decomposition of the prospective input space. More broadly, while extensive integrity checks can confirm database resolvability and linkage consistency, it is not feasible to validate the correctness of every transformed exchange across a large scenario-modified inventory. The approach manages this through strict versioning/freezing, bounded and logged overrides at hotspots, and targeted sensitivities rather than exhaustive verification. Finally, the system boundary emphasized EoL modules C3–C4 and Stage D to isolate EoL pathway and substitution effects; adoption and market penetration dynamics, design evolution, and region-specific policy feedback were not explicitly modelled and remain important extensions.

5.4 Future Work

The framework suggests several high-value extensions that directly follow from the thesis objectives. On the background side, expanding sectoral transformation beyond energy and fuels and improving Canada-specific spatial fidelity would strengthen interpretation of both foreground burdens and displaced-system intensities. premise also supports integration of custom scenarios in addition to standard IAM pathways, offering a clear route to future work that constructs Canada-tailored futures (for example, national net-zero, current-policy, and global net-zero alignment for electricity and fuels) while retaining consistent foreground EoL models and Stage D governance. On the uncertainty side, enabling defensible propagation of background exchange uncertainty in IAM-coupled inventories, and thereby implementing prospective joint (background + foreground) uncertainty attribution, remains an important methodological frontier. On the consequential interpretation side, targeted structural stress tests of substitution governance (market boundaries, marginal response representations, and equivalence ratios) would complement parameter uncertainty and better map persistence conditions for net claims in credit-dominant regimes. For emerging technologies, empirical validation of key scale-up parameters, eligibility/quality-gate behaviour, and scrap-quality constraints would reduce reliance on bounded priors and improve decision relevance. Extending the case study beyond aluminium to additional façade materials and incorporating deeper non-climate mechanism interpretation would further generalize insights about circular EoL strategies under prospective system evolution.

References

Al-Alimi, S., Yusuf, N. K., Ghaleb, A. M., Lajis, M. A., Shamsudin, S., Zhou, W., Altharan, Y. M., Abdulwahab, H. S., Saif, Y., Didane, D. H., S t t, I., & Adam, A. (2024). Recycling aluminium for sustainable development: A review of different processing technologies in green manufacturing. *Results in Engineering*, *23*, 102566. https://doi.org/10.1016/j.rineng.2024.102566

Aluminium Association of Canada. (2025). Over 100 years of history. *Aluminium.Ca*. https://aluminium.ca/en/the-industry/over-100-years-of-history/

Amor, M. B., Gaudreault, C., Pineau, P.-O., & Samson, R. (2014). Implications of integrating electricity supply dynamics into life cycle assessment: A case study of renewable distributed generation. *Renewable Energy*, *69*, 410–419. https://doi.org/10.1016/j.renene.2014.03.063

Anand, C. K., & Amor, B. (2017). Recent developments, future challenges and new research directions in LCA of buildings: A critical review. *Renewable and Sustainable Energy Reviews*, *67*, 408–416. https://doi.org/10.1016/j.rser.2016.09.058

Arvidsson, R., Nguyen, D., & Svanström, M. (2015). Life Cycle Assessment of Cellulose Nanofibrils Production by Mechanical Treatment and Two Different Pretreatment Processes. *Environmental Science & Technology*, *49*(11), 6881–6890. https://doi.org/10.1021/acs.est.5b00888

Arvidsson, R., Sandén, B. A., & Svanström, M. (2023). *Prospective, Anticipatory and Ex-Ante – What's the Difference? Sorting Out Concepts for Time-Related LCA*.

Arvidsson, R., Tillman, A., Sandén, B. A., Janssen, M., Nordelöf, A., Kushnir, D., & Molander, S. (2018). Environmental Assessment of Emerging Technologies: Recommendations for Prospective LCA. *Journal of Industrial Ecology*, *22*(6), 1286–1294. https://doi.org/10.1111/jiec.12690

Ayres, R. U. (1995). Life cycle analysis: A critique. *Resources, Conservation and Recycling, Life Cycle Management*, *14*(3), 199–223. https://doi.org/10.1016/0921-3449(95)00017-D

Baker, J., & Lepech, M. (2009). Treatment of uncertainties in life cycle assessment. *10th International Conference on Structural Safety and Reliability (ICOSSAR09), Osaka, Japan, (2009)*.

Bamber, N., Turner, I., Arulnathan, V., Li, Y., Zargar Ershadi, S., Smart, A., & Pelletier, N. (2020). Comparing sources and analysis of uncertainty in consequential and attributional life cycle assessment: Review of current practice and recommendations. *The International Journal of Life Cycle Assessment*, *25*(1), 168–180. https://doi.org/10.1007/s11367-019-01663-1

Bare, J. C. (2010). Life cycle impact assessment research developments and needs. *Clean Technologies and Environmental Policy*, *12*(4), 341–351. https://doi.org/10.1007/s10098-009-0265-9

Beloin-Saint-Pierre, D., Albers, A., Hélias, A., Tiruta-Barna, L., Fantke, P., Levasseur, A., Benetto, E., Benoist, A., & Collet, P. (2020). Addressing temporal considerations in life cycle assessment. *Science of The Total Environment*, *743*, 140700. https://doi.org/10.1016/j.scitotenv.2020.140700

Beloin-Saint-Pierre, D., Heijungs, R., & Blanc, I. (2014). The ESPA (Enhanced Structural Path Analysis) method: A solution to an implementation challenge for dynamic life cycle assessment studies. *The International Journal of Life Cycle Assessment*, *19*(4), 861–871. https://doi.org/10.1007/s11367-014-0710-9

Bergerson, J. A., Brandt, A., Cresko, J., Carbajales-Dale, M., MacLean, H. L., Matthews, H. S., McCoy, S., McManus, M., Miller, S. A., Morrow, W. R., Posen, I. D., Seager, T., Skone, T., & Sleep, S. (2020). Life cycle assessment of emerging technologies: Evaluation techniques at different stages of market and technical maturity. *Journal of Industrial Ecology*, *24*(1), 11–25. https://doi.org/10.1111/jiec.12954

Beylot, A. (2024). *Advancement of the LCA method to account for the dissipation of mineral resources and associated damage over the life cycle of products and systems: Validation of Acquired Experience (VAE)*. Chemical and Process Engineering. Université de Bordeaux.

Bicalho, T., Sauer, I., Rambaud, A., & Altukhova, Y. (2017). LCA data quality: A management science perspective. *Journal of Cleaner Production*, *156*, 888–898. https://doi.org/10.1016/j.jclepro.2017.03.229

Björklund, A. E. (2002). Survey of approaches to improve reliability in lca. *The International Journal of Life Cycle Assessment*, *7*(2), 64–72. https://doi.org/10.1007/BF02978849

Boin, U. M. J., & Bertram, M. (2005). Melting standardized aluminum scrap: A mass balance model for europe. *JOM*, *57*(8), 26–33. https://doi.org/10.1007/s11837-005-0164-4

Brandão, M., Weidema, B. P., Martin, M., Cowie, A., Hamelin, L., & Zamagni, A. (2024). Consequential Life Cycle Assessment: What, Why and How? In *Encyclopedia of Sustainable Technologies* (pp. 181–189). Elsevier. https://doi.org/10.1016/B978-0-323-90386-8.00001-2

BSI. (2026). *BS en 15978 | PDF*. https://www.scribd.com/document/839251653/BS-EN-15978

Buyle, M., Audenaert, A., Billen, P., Boonen, K., & Van Passel, S. (2019). The Future of Ex-Ante LCA? Lessons Learned and Practical Recommendations. *Sustainability*, *11*(19), 5456. https://doi.org/10.3390/su11195456

Buyle, M., Galle, W., Debacker, W., & Audenaert, A. (2019). Sustainability assessment of circular building alternatives: Consequential LCA and LCC for internal wall assemblies as a case study in a Belgian context. *Journal of Cleaner Production*, *218*, 141–156. https://doi.org/10.1016/j.jclepro.2019.01.306

Caduff, M., Huijbregts, M. A. J., Althaus, H.-J., Koehler, A., & Hellweg, S. (2012). Wind Power Electricity: The Bigger the Turbine, The Greener the Electricity? *Environmental Science & Technology*, *46*(9), 4725–4733. https://doi.org/10.1021/es204108n

Caduff, M., Huijbregts, M. A. J., Koehler, A., Althaus, H., & Hellweg, S. (2014). Scaling Relationships in Life Cycle Assessment: The Case of Heat Production from Biomass and Heat Pumps. *Journal of Industrial Ecology*, *18*(3), 393–406. https://doi.org/10.1111/jiec.12122

Cherubini, E., Franco, D., Zanghelini, G. M., & Soares, S. R. (2018). Uncertainty in LCA case study due to allocation approaches and life cycle impact assessment methods. *The International Journal of Life Cycle Assessment*, *23*(10), 2055–2070. https://doi.org/10.1007/s11367-017-1432-6

Cherubini, F., Peters, G. P., Berntsen, T., Strømman, A. H., & Hertwich, E. (2011). CO2 emissions from biomass combustion for bioenergy: Atmospheric decay and contribution to global warming: GLOBAL WARMING POTENTIAL OF CO2 FROM BIOENERGY. *GCB Bioenergy*, *3*(5), 413–426. https://doi.org/10.1111/j.1757-1707.2011.01102.x

Codotto, G., Pizzol, M., Hegland, T. J., & Madsen, N. (2024). Model uncertainty versus variability in the life cycle assessment of commercial fisheries. *Journal of Industrial Ecology*, *28*(1), 160–172. https://doi.org/10.1111/jiec.13453

Collingridge, D. (1980). *The social control of technology*. https://scholar.google.com/scholar_lookup?title=The%20Social%20Control%20of%20Technology&publication_year=1980&author=D.%20Collingridge

Consequential-LCA. (2020). *Marginal Suppliers*. https://consequential-lca.org/clca/marginal-suppliers/

Cooper, D. (2013). *Reuse of steel and aluminium without melting*.

Cooper, D. R., & Allwood, J. M. (2012). Reusing Steel and Aluminum Components at End of Product Life. *Environmental Science & Technology*, *46*(18), 10334–10340. https://doi.org/10.1021/es301093a

Cooper, D. R., Song, J., & Gerard, R. (2018). Metal recovery during melting of extruded machining chips. *Journal of Cleaner Production*, *200*, 282–292. https://doi.org/10.1016/j.jclepro.2018.07.246

Cox, B., Bauer, C., Mendoza Beltran, A., Van Vuuren, D. P., & Mutel, C. L. (2020). Life cycle environmental and cost comparison of current and future passenger cars under different energy scenarios. *Applied Energy*, *269*, 115021. https://doi.org/10.1016/j.apenergy.2020.115021

Cox, B., Mutel, C. L., Bauer, C., Mendoza Beltran, A., & Van Vuuren, D. P. (2018). Uncertain Environmental Footprint of Current and Future Battery Electric Vehicles. *Environmental Science & Technology*, *52*(8), 4989–4995. https://doi.org/10.1021/acs.est.8b00261

Cucurachi, S., Van Der Giesen, C., & Guinée, J. (2018). Ex-ante LCA of Emerging Technologies. *Procedia CIRP*, *69*, 463–468. https://doi.org/10.1016/j.procir.2017.11.005

Curran, M. A., Mann, M., & Norris, G. (2005). The international workshop on electricity data for life cycle inventories. *Journal of Cleaner Production*, *13*(8), 853–862. https://doi.org/10.1016/j.jclepro.2002.03.001

Deebika, P., & Saravanakumar, M. P. (2023). Utilization of RO rejects and waste aluminum scraps for hydrogen generation. *International Journal of Hydrogen Energy*, *48*(91), 35516–35531. https://doi.org/10.1016/j.ijhydene.2023.06.015

Demirci, I., Isselin, J., & Créquy, S. (2025). Mechanical performance of bulk products obtained by solid state recycling of aluminium machining chips. *The International Journal of Advanced Manufacturing Technology*, *137*(11–12), 5499–5510. https://doi.org/10.1007/s00170-025-15498-8

DOE. (2023). *Prototype Building Models | Building Energy Codes Program*. https://www.energycodes.gov/prototype-building-models

Duflou, J. R., Tekkaya, A. E., Haase, M., Welo, T., Vanmeensel, K., Kellens, K., Dewulf, W., & Paraskevas, D. (2015). Environmental assessment of solid state recycling routes for aluminium alloys: Can solid state processes significantly reduce the environmental impact of aluminium recycling? *CIRP Annals*, *64*(1), 37–40. https://doi.org/10.1016/j.cirp.2015.04.051

Earles, J. M., & Halog, A. (2011). Consequential life cycle assessment: A review. *The International Journal of Life Cycle Assessment*, *16*(5), 445–453. https://doi.org/10.1007/s11367-011-0275-9

ecoinvent Association. (2025). *Market Activities*. https://support.ecoinvent.org/market-activities

Ekvall, T. (2020). Attributional and Consequential Life Cycle Assessment. In M. José Bastante-Ceca, J. Luis Fuentes-Bargues, L. Hufnagel, F.-C. Mihai, & C. Iatu (Eds.), *Sustainability Assessment at the 21st century*. IntechOpen. https://doi.org/10.5772/intechopen.89202

Ekvall, T., Tillman, A.-M., & Molander, S. (2005). Normative ethics and methodology for life cycle assessment. *Journal of Cleaner Production*, *13*(13–14), 1225–1234. https://doi.org/10.1016/j.jclepro.2005.05.010

Ekvall, T., & Weidema, B. P. (2004). System boundaries and input data in consequential life cycle inventory analysis. *The International Journal of Life Cycle Assessment*, *9*(3), 161–171. https://doi.org/10.1007/BF02994190

Fava, J. A., Jensen, A. A., Lindfors, L., Pomper, S., De Smet, B., Warren, J., & Vigon, B. (1994). *Life-cycle assessment data quality: A conceptual framework*. Society of Environmental Toxicology and Chemistry and SETAC Foundation for ….

Finnveden, G. (2000). On the limitations of life cycle assessment and environmental systems analysis tools in general. *The International Journal of Life Cycle Assessment*, *5*(4), 229. https://doi.org/10.1007/BF02979365

Finnveden, G., Hauschild, M. Z., Ekvall, T., Guinée, J., Heijungs, R., Hellweg, S., Koehler, A., Pennington, D., & Suh, S. (2009). Recent developments in Life Cycle Assessment. *Journal of Environmental Management*, *91*(1), 1–21. https://doi.org/10.1016/j.jenvman.2009.06.018

Fouquet, M., Levasseur, A., Margni, M., Lebert, A., Lasvaux, S., Souyri, B., Buhé, C., & Woloszyn, M. (2015). Methodological challenges and developments in LCA of low energy buildings: Application to biogenic carbon and global warming assessment. *Building and Environment*, *90*, 51–59. https://doi.org/10.1016/j.buildenv.2015.03.022

Frischknecht, R., Althaus, H.-J., Doka, G., Dones, R., Heck, T., Hellweg, S., Hischier, R., Nemecek, T., Rebitzer, G., Spielmann, M., & Wernet, G. (2007). *Data v2.0 (2007) Rolf Frischknecht, Niels Jungbluth (Editors)*. (1).

Gavankar, S., Suh, S., & Keller, A. A. (2015). The Role of Scale and Technology Maturity in Life Cycle Assessment of Emerging Technologies: A Case Study on Carbon Nanotubes. *Journal of Industrial Ecology*, *19*(1), 51–60. https://doi.org/10.1111/jiec.12175

Georgiades, M., Shah, I. H., Steubing, B., Cheeseman, C., & Myers, R. J. (2023). Prospective life cycle assessment of European cement production. *Resources, Conservation and Recycling*, *194*, 106998. https://doi.org/10.1016/j.resconrec.2023.106998

Gibon, T., Wood, R., Arvesen, A., Bergesen, J. D., Suh, S., & Hertwich, E. G. (2015). A Methodology for Integrated, Multiregional Life Cycle Assessment Scenarios under Large-Scale Technological Change. *Environmental Science & Technology*, *49*(18), 11218–11226. https://doi.org/10.1021/acs.est.5b01558

Global Affairs Canada. (2020, December 16). *Aluminium*. GAC. https://www.international.gc.ca/trade-commerce/controls-controles/aluminum-aluminium/index.aspx?lang=eng

Government of Canada, C. E. R. (2023, November 24). *CER – Scenarios and Assumptions*. https://www.cer-rec.gc.ca/en/data-analysis/canada-energy-future/2023/scenarios-and-assumptions/

Groen, E. A., Heijungs, R., Bokkers, E. A., & De Boer, I. J. (2014a). Sensitivity analysis in life cycle assessment. *Proceedings of the 9th International Conference on Life Cycle Assessment in the Agri-Food Sector*, 8–10. https://www.researchgate.net/profile/Evelyne-Groen/publication/283417261_Sensitivity_analysis_in_life_cycle_assessment/links/5a7da50caca272341aef9542/Sensitivity-analysis-in-life-cycle-assessment.pdf

Groen, E. A., Heijungs, R., Bokkers, E. A. M., & De Boer, I. J. M. (2014b). Methods for uncertainty propagation in life cycle assessment. *Environmental Modelling & Software*, *62*, 316–325. https://doi.org/10.1016/j.envsoft.2014.10.006

Guinée, J. B., Cucurachi, S., Henriksson, P. J. G., & Heijungs, R. (2018). Digesting the alphabet soup of LCA. *The International Journal of Life Cycle Assessment*, *23*(7), 1507–1511. https://doi.org/10.1007/s11367-018-1478-0

Guo, M., & Murphy, R. J. (2012). LCA data quality: Sensitivity and uncertainty analysis. *Science of The Total Environment*, *435–436*, 230–243. https://doi.org/10.1016/j.scitotenv.2012.07.006

Gütschow, J., Jeffery, M. L., Günther, A., & Meinshausen, M. (2021). Country-resolved combined emission and socio-economic pathways based on the Representative Concentration Pathway (RCP) and Shared Socio-Economic Pathway (SSP) scenarios. *Earth System Science Data*, *13*(3), 1005–1040. https://doi.org/10.5194/essd-13-1005-2021

Harmsen, M., Kriegler, E., Van Vuuren, D. P., Van Der Wijst, K.-I., Luderer, G., Cui, R., Dessens, O., Drouet, L., Emmerling, J., Morris, J. F., Fosse, F., Fragkiadakis, D., Fragkiadakis, K., Fragkos, P., Fricko, O., Fujimori, S., Gernaat, D., Guivarch, C., Iyer, G., … Zakeri, B. (2021). Integrated assessment model diagnostics: Key indicators and model evolution. *Environmental Research Letters*, *16*(5), 054046. https://doi.org/10.1088/1748-9326/abf964

Hauschild, M. Z., Rosenbaum, R. K., & Olsen, S. I. (Eds.). (2018). *Life Cycle Assessment: Theory and Practice*. Springer International Publishing. https://doi.org/10.1007/978-3-319-56475-3

Heeren, N., & Hellweg, S. (2019). Tracking Construction Material over Space and Time: Prospective and Geo-referenced Modeling of Building Stocks and Construction Material Flows. *Journal of Industrial Ecology*, *23*(1), 253–267. https://doi.org/10.1111/jiec.12739

Heijungs, R., & Huijbregts, M. A. J. (2004). *A Review of Approaches to Treat Uncertainty in LCA*.

Hellweg, S., Hofstetter, T. B., & Hungerbuhler, K. (2003). Discounting and the environment should current impacts be weighted differently than impacts harming future generations? *The International Journal of Life Cycle Assessment*, *8*(1), 8. https://doi.org/10.1007/BF02978744

Hertwich, E. G., Gibon, T., Bouman, E. A., Arvesen, A., Suh, S., Heath, G. A., Bergesen, J. D., Ramirez, A., Vega, M. I., & Shi, L. (2015). Integrated life-cycle assessment of electricity-supply scenarios confirms global environmental benefit of low-carbon technologies. *Proceedings of the National Academy of Sciences*, *112*(20), 6277–6282. https://doi.org/10.1073/pnas.1312753111

Hetherington, A. C., Borrion, A. L., Griffiths, O. G., & McManus, M. C. (2014). Use of LCA as a development tool within early research: Challenges and issues across different sectors. *The*

International Journal of Life Cycle Assessment, *19*(1), 130–143. https://doi.org/10.1007/s11367-013-0627-8

Hoeting, J. A., Madigan, D., Raftery, A. E., & Volinsky, C. T. (1999). *Bayesian Model Averaging: A Tutorial*.

Huijbregts, M. A. J. (1998a). Application of uncertainty and variability in LCA. *The International Journal of Life Cycle Assessment*, *3*(5), 273. https://doi.org/10.1007/BF02979835

Huijbregts, M. A. J. (1998b). Part II: Dealing with parameter uncertainty and uncertainty due to choices in life cycle assessment. *The International Journal of Life Cycle Assessment*, *3*(6), 343–351. https://doi.org/10.1007/BF02979345

Huijbregts, M. A. J., Gilijamse, W., Ragas, A. M. J., & Reijnders, L. (2003). Evaluating Uncertainty in Environmental Life-Cycle Assessment. A Case Study Comparing Two Insulation Options for a Dutch One-Family Dwelling. *Environmental Science & Technology*, *37*(11), 2600–2608. https://doi.org/10.1021/es020971+

Igos, E., Benetto, E., Meyer, R., Baustert, P., & Othoniel, B. (2019). How to treat uncertainties in life cycle assessment studies? *The International Journal of Life Cycle Assessment*, *24*(4), 794–807. https://doi.org/10.1007/s11367-018-1477-1

Ingarao, G., Amato, M., Latif, A., La Rosa, A. D., Di Lorenzo, R., & Fratini, L. (2023). Life Cycle Assessment of aluminum alloys chips recycling through single and multi-step Friction Stir Consolidation processes. *Journal of Manufacturing Systems*, *68*, 651–659. https://doi.org/10.1016/j.jmsy.2023.05.021

Intergovernmental Panel On Climate Change (IPCC) (Ed.). (2023). Annex III: Scenarios and Modelling Methods. In *Climate Change 2022—Mitigation of Climate Change* (1st ed., pp. 1841–1908). Cambridge University Press. https://doi.org/10.1017/9781009157926.022

ISO. (2006a). *ISO 14040:2006—Environmental management—Life cycle assessment—Principles and framework*. https://www.iso.org/standard/37456.html

ISO. (2006b). *ISO 14044:2006*. ISO. https://www.iso.org/standard/38498.html

Jones, C., Gilbert, P., Raugei, M., Mander, S., & Leccisi, E. (2017). An approach to prospective consequential life cycle assessment and net energy analysis of distributed electricity generation. *Energy Policy*, *100*, 350–358. https://doi.org/10.1016/j.enpol.2016.08.030

Kendall, A., & Price, L. (2012). Incorporating Time-Corrected Life Cycle Greenhouse Gas Emissions in Vehicle Regulations. *Environmental Science & Technology*, *46*(5), 2557–2563. https://doi.org/10.1021/es203098j

Kennedy, D. J., Montgomery, D. C., & Quay, B. H. (1996). Data quality. *The International Journal of Life Cycle Assessment*, *1*(4), 199–207. https://doi.org/10.1007/BF02978693

Khoo, H. H., Tan, R. B. H., & Chng, K. W. L. (2010). Environmental impacts of conventional plastic and bio-based carrier bags: Part 1: Life cycle production. *The International Journal of Life Cycle Assessment*, *15*(3), 284–293. https://doi.org/10.1007/s11367-010-0162-9

Kravchenko, M., Pigosso, D. Ca., & McAloone, T. C. (2019). Towards the ex-ante sustainability screening of circular economy initiatives in manufacturing companies: Consolidation of leading sustainability-related performance indicators. *Journal of Cleaner Production*, *241*, 118318. https://doi.org/10.1016/j.jclepro.2019.118318

Kriegler, E., O'Neill, B. C., Hallegatte, S., Kram, T., Lempert, R. J., Moss, R. H., & Wilbanks, T. (2012). The need for and use of socio-economic scenarios for climate change analysis: A new approach based on shared socio-economic pathways. *Global Environmental Change*, *22*(4), 807–822. https://doi.org/10.1016/j.gloenvcha.2012.05.005

Lacirignola, M. (2017). *Life cycle assessment of enhanced geothermal systems: From specific case studies to generic parameterized models* [PhD Thesis, Conservatoire national des arts et metiers-CNAM]. https://theses.hal.science/tel-01591097/

Latif, A., Ingarao, G., & Fratini, L. (2022). Multi-material based functionally graded billets manufacturing through friction stir consolidation of aluminium alloys chips. *CIRP Annals*, *71*(1), 261–264. https://doi.org/10.1016/j.cirp.2022.03.035

Latif, A., Ingarao, G., Gucciardi, M., & Fratini, L. (2022). A novel approach to enhance mechanical properties during recycling of aluminum alloy scrap through friction stir consolidation. *The International Journal of Advanced Manufacturing Technology*, *119*(3–4), 1989–2005. https://doi.org/10.1007/s00170-021-08346-y

Lebailly, F., Levasseur, A., Samson, R., & Deschênes, L. (2014). Development of a dynamic LCA approach for the freshwater ecotoxicity impact of metals and application to a case study regarding zinc fertilization. *The International Journal of Life Cycle Assessment*, *19*(10), 1745–1754. https://doi.org/10.1007/s11367-014-0779-1

Levasseur, A., Lesage, P., Margni, M., Deschênes, L., & Samson, R. (2010). Considering Time in LCA: Dynamic LCA and Its Application to Global Warming Impact Assessments. *Environmental Science & Technology*, *44*(8), 3169–3174. https://doi.org/10.1021/es9030003

Lloyd, S. M., & Ries, R. (2007). Characterizing, Propagating, and Analyzing Uncertainty in Life-Cycle Assessment: A Survey of Quantitative Approaches. *Journal of Industrial Ecology*, *11*(1), 161–179. https://doi.org/10.1162/jiec.2007.1136

Luderer, G., Leimbach, M., Bauer, N., Kriegler, E., Baumstark, L., Giannousakis, A., Hilaire, J., Klein, D., Levesque, A., Pehl, M., Pietzcker, R., Piontek, F., Roming, N., Schwanitz, V. J., & Strefler, J. (2015). *Description of the REMIND model (Version 1.6)*.

Lueddeckens, S., Saling, P., & Guenther, E. (2020). Temporal issues in life cycle assessment—A systematic review. *The International Journal of Life Cycle Assessment*, *25*(8), 1385–1401. https://doi.org/10.1007/s11367-020-01757-1

Maes, B., Audenaert, A., Craeye, B., & Buyle, M. (2021). Consequential ex-ante life cycle assessment on clinker production in the EU: How will the future influence its environmental impact? *Journal of Cleaner Production*, *315*, 128081. https://doi.org/10.1016/j.jclepro.2021.128081

Maes, B., Craeye, B., Buyle, M., & Audenaert, A. (2024). Enriching IAM scenarios for effective pLCA integration: A clinker case study. *Journal of Cleaner Production*, *470*, 143316. https://doi.org/10.1016/j.jclepro.2024.143316

Maes, B., Sacchi, R., Steubing, B., Pizzol, M., Audenaert, A., Craeye, B., & Buyle, M. (2023). Prospective consequential life cycle assessment: Identifying the future marginal suppliers using integrated assessment models. *Renewable and Sustainable Energy Reviews*, *188*, 113830. https://doi.org/10.1016/j.rser.2023.113830

Marsh, E., Allen, S., & Hattam, L. (2023). Tackling uncertainty in life cycle assessments for the built environment: A review. *Building and Environment*, *231*, 109941. https://doi.org/10.1016/j.buildenv.2022.109941

Maurice, B., Frischknecht, R., Coelho-Schwirtz, V., & Hungerbuhler, K. (2000). Uncertainty analysis in life cycle inventory. Application to the production of electricity with French coal power plants. *Journal of Cleaner Production*.

MD, M. (1991). Morris Factorial Sampling Plans for Preliminary Computational Experiments. *Technometrics*, *33*, 161. https://doi.org/10.2307/1269043

Mendoza Beltran, A., Cox, B., Mutel, C., Van Vuuren, D. P., Font Vivanco, D., Deetman, S., Edelenbosch, O. Y., Guinée, J., & Tukker, A. (2020). When the Background Matters: Using Scenarios from Integrated Assessment Models in Prospective Life Cycle Assessment. *Journal of Industrial Ecology*, *24*(1), 64–79. https://doi.org/10.1111/jiec.12825

Milà i Canals, L., Azapagic, A., Doka, G., Jefferies, D., King, H., Mutel, C. L., Nemecek, T., Roches, A., Sim, S., Stichnothe, H., Thoma, G., & Williams, A. (2011). *Approaches for Addressing Life Cycle Assessment Data Gaps for Bio-based Products*. https://onlinelibrary.wiley.com/doi/full/10.1111/j.1530-9290.2011.00369.x?casa_token=XKazv5udyUIAAAAA%3A2YpOZYf9_lyurvFC7bFUYI5ATHRs0-iTJNb8ViMblmlxGQiglqHp8yU6iGi_29ojRQQVPasxCAN-_Pk

Moni, S. M., Mahmud, R., High, K., & Carbajales-Dale, M. (2020). Life cycle assessment of emerging technologies: A review. *Journal of Industrial Ecology*, *24*(1), 52–63. https://doi.org/10.1111/jiec.12965

Moss, R. H., Edmonds, J. A., Hibbard, K. A., Manning, M. R., Rose, S. K., Van Vuuren, D. P., Carter, T. R., Emori, S., Kainuma, M., Kram, T., Meehl, G. A., Mitchell, J. F. B., Nakicenovic, N., Riahi, K., Smith, S. J., Stouffer, R. J., Thomson, A. M., Weyant, J. P., & Wilbanks, T. J. (2010). The next generation of scenarios for climate change research and assessment. *Nature*, *463*(7282), 747–756. https://doi.org/10.1038/nature08823

Muñoz, I., & Weidema, B. (2023). Example—Marginal Electricity in Denmark. Version: 2023-09-20. *Consequential LCA*. www.consequential-lca.org

Mutel, C. L., De Baan, L., & Hellweg, S. (2013). Two-Step Sensitivity Testing of Parametrized and Regionalized Life Cycle Assessments: Methodology and Case Study. *Environmental Science & Technology*, *47*(11), 5660–5667. https://doi.org/10.1021/es3050949

Mutel, C., Liao, X., Patouillard, L., Bare, J., Fantke, P., Frischknecht, R., Hauschild, M., Jolliet, O., Maia De Souza, D., Laurent, A., Pfister, S., & Verones, F. (2019). Overview and recommendations for regionalized life cycle impact assessment. *The International Journal of Life Cycle Assessment*, *24*(5), 856–865. https://doi.org/10.1007/s11367-018-1539-4

NRCan. (2018, January 19). *Aluminum Facts*. https://natural-resources.canada.ca/minerals-mining/mining-data-statistics-analysis/minerals-metals-facts/aluminum-facts

NRCan. (2024, February 13). *Hydrogen Strategy for Canada: Progress Report*. https://natural-resources.canada.ca/energy-sources/clean-fuels/hydrogen-strategy/hydrogen-strategy-canada-progress-report

Nurdiawati, A., Mir, B. A., & Al-Ghamdi, S. G. (2025). Recent advancements in prospective life cycle assessment: Current practices, trends, and implications for future research. *Resources, Environment and Sustainability*, *20*, 100203. https://doi.org/10.1016/j.resenv.2025.100203

O'Neill, B. C., Kriegler, E., Riahi, K., Ebi, K. L., Hallegatte, S., Carter, T. R., Mathur, R., & Van Vuuren, D. P. (2014). A new scenario framework for climate change research: The concept of shared socioeconomic pathways. *Climatic Change*, *122*(3), 387–400. https://doi.org/10.1007/s10584-013-0905-2

Pehnt, M. (2006). Dynamic life cycle assessment (LCA) of renewable energy technologies. *Renewable Energy*, *31*(1), 55–71. https://doi.org/10.1016/j.renene.2005.03.002

Peñaloza, D., Erlandsson, M., & Falk, A. (2016). Exploring the climate impact effects of increased use of bio-based materials in buildings. *Construction and Building Materials*, *125*, 219–226. https://doi.org/10.1016/j.conbuildmat.2016.08.041

Peters, J. F., Baumann, M., Zimmermann, B., Braun, J., & Weil, M. (2017). The environmental impact of Li-Ion batteries and the role of key parameters – A review. *Renewable and Sustainable Energy Reviews*, *67*, 491–506. https://doi.org/10.1016/j.rser.2016.08.039

Petrovic, J., & Thomas, G. (2011). *Reaction of Aluminum with Water to Produce Hydrogen: A Study of Issues Related to the Use of Aluminum for On-Board Vehicular Hydrogen Storage (Version 2 - 2010 Update)*. US Department of Energy (USDOE), Washington, DC (United States). Office of Energy Efficiency and Renewable Energy (EERE), Office of Sustainable Transportation. Hydrogen Fuel Cell Technologies Office (HFTO). https://doi.org/10.2172/1219359

RAICO. (2026). *RAICO - THERM+ H-I – Timber façade*. https://www.raico.de/en/products/facades-content/timber/therm-h-i.html

Ravikumar, D., Seager, T. P., Cucurachi, S., Prado, V., & Mutel, C. (2018). Novel Method of Sensitivity Analysis Improves the Prioritization of Research in Anticipatory Life Cycle Assessment of Emerging Technologies. *Environmental Science & Technology*, *52*(11), 6534–6543. https://doi.org/10.1021/acs.est.7b04517

Reap, J., Roman, F., Duncan, S., & Bras, B. (2008). A survey of unresolved problems in life cycle assessment. *The International Journal of Life Cycle Assessment*, *13*(4), 290–300. https://doi.org/10.1007/s11367-008-0008-x

Riahi, K., Rao, S., Krey, V., Cho, C., Chirkov, V., Fischer, G., Kindermann, G., Nakicenovic, N., & Rafaj, P. (2011). RCP 8.5—A scenario of comparatively high greenhouse gas emissions. *Climatic Change*, *109*(1–2), 33–57. https://doi.org/10.1007/s10584-011-0149-y

Riahi, K., Van Vuuren, D. P., Kriegler, E., Edmonds, J., O'Neill, B. C., Fujimori, S., Bauer, N., Calvin, K., Dellink, R., Fricko, O., Lutz, W., Popp, A., Cuaresma, J. C., Kc, S., Leimbach, M., Jiang, L., Kram, T., Rao, S., Emmerling, J., … Tavoni, M. (2017). The Shared Socioeconomic Pathways and their energy, land use, and greenhouse gas emissions implications: An overview. *Global Environmental Change*, *42*, 153–168. https://doi.org/10.1016/j.gloenvcha.2016.05.009

Rio Tinto. (2023). *Rio Tinto to expand its AP60 low-carbon aluminium smelter in Quebec*. https://www.riotinto.com/fr-ca/can/news/releases/2023/rio-tinto-to-expand-its-ap60-low-carbon-aluminium-smelter-in-quebec

Sacchi, R., Terlouw, T., Siala, K., Dirnaichner, A., Bauer, C., Cox, B., Mutel, C., Daioglou, V., & Luderer, G. (2022). PRospective EnvironMental Impact asSEment (premise): A streamlined approach to producing databases for prospective life cycle assessment using integrated assessment models. *Renewable and Sustainable Energy Reviews*, *160*, 112311. https://doi.org/10.1016/j.rser.2022.112311

Sacchi, R., Terlouw, T., Siala, K., Dirnaichner, A., Bauer, C., Cox, B., Mutel, C., Daioglou, V., & Luderer, G. (2023). *Premise* (Version 1.4.2) [Jupyter Notebook]. https://github.com/polca/premise (Original work published 2018)

Sandén, B. A., & Karlström, M. (2007). Positive and negative feedback in consequential life-cycle assessment. *Journal of Cleaner Production*, *15*(15), 1469–1481. https://doi.org/10.1016/j.jclepro.2006.03.005

Schaubroeck, T. (2023). Relevance of attributional and consequential life cycle assessment for society and decision support. *Frontiers in Sustainability*, *4*, 1063583. https://doi.org/10.3389/frsus.2023.1063583

Schrijvers, D., Loubet, P., & Sonnemann, G. (2020). Archetypes of Goal and Scope Definitions for Consistent Allocation in LCA. *Sustainability*, *12*(14), 5587. https://doi.org/10.3390/su12145587

Shen, L., Nieuwlaar, E., Worrell, E., & Patel, M. K. (2011). Life cycle energy and GHG emissions of PET recycling: Change-oriented effects. *The International Journal of Life Cycle Assessment*, *16*(6), 522–536. https://doi.org/10.1007/s11367-011-0296-4

Shimako, A. H., Tiruta-Barna, L., & Ahmadi, A. (2017). Operational integration of time dependent toxicity impact category in dynamic LCA. *Science of The Total Environment*, *599–600*, 806–819. https://doi.org/10.1016/j.scitotenv.2017.04.211

Shimako, A. H., Tiruta-Barna, L., Bisinella De Faria, A. B., Ahmadi, A., & Spérandio, M. (2018). Sensitivity analysis of temporal parameters in a dynamic LCA framework. *Science of The Total Environment*, *624*, 1250–1262. https://doi.org/10.1016/j.scitotenv.2017.12.220

Šimaitis, J., Butnar, I., Sacchi, R., Lupton, R., Vagg, C., & Allen, S. (2025). Expanding scenario diversity in prospective LCA: Coupling the TIAM-UCL integrated assessment model with Premise and ecoinvent. *Renewable and Sustainable Energy Reviews*, *211*, 115298. https://doi.org/10.1016/j.rser.2024.115298

Singhofen, A., Hemming, C. R., Weidema, B. P., Grisel, L., Bretz, R., De Smet, B., & Russell, D. (1996). Life cycle inventory data: Development of a common format. *The International Journal of Life Cycle Assessment*, *1*(3), 171–178. https://doi.org/10.1007/BF02978948

Steen, B. (1997). On uncertainty and sensitivity of LCA-based priority setting. *Journal of Cleaner Production*, *5*(4), 255–262. https://doi.org/10.1016/S0959-6526(97)00039-5

Stehfest, E., van Vuuren, D., Kram, T., Bouwman, L., Alkemade, R., Bakkenes, M., Bouwman, A., den Elzen, M., Janse, M., Lucas, P., van Minnen, J., Prins, A. G., Biemans, H., & Müller, C. (2014). *Integrated Assessment of Global Environmental Change with IMAGE 3.0 (Report No. 735)*. PBL Netherlands Environmental Assessment Agency. https://www.pbl.nl/en/publications/integrated-assessment-of-global-environmental-change-with-image-3-0

Steubing, B., Wernet, G., Reinhard, J., Bauer, C., & Moreno-Ruiz, E. (2016). The ecoinvent database version 3 (part II): Analyzing LCA results and comparison to version 2. *The*

International Journal of Life Cycle Assessment, *21*(9), 1269–1281. https://doi.org/10.1007/s11367-016-1109-6

Tanguay, X., Essoua Essoua, G. G., & Amor, B. (2021). Attributional and consequential life cycle assessments in a circular economy with integration of a quality indicator: A case study of cascading wood products. *Journal of Industrial Ecology*, *25*(6), 1462–1473. https://doi.org/10.1111/jiec.13167

Thonemann, N., Schulte, A., & Maga, D. (2020). How to Conduct Prospective Life Cycle Assessment for Emerging Technologies? A Systematic Review and Methodological Guidance. *Sustainability*, *12*(3), 1192. https://doi.org/10.3390/su12031192

Tillman, A.-M. (2000). Significance of decision-making for LCA methodology. *Environmental Impact Assessment Review*, *20*(1), 113–123. https://doi.org/10.1016/S0195-9255(99)00035-9

Tiruta-Barna, L., Pigné, Y., Navarrete Gutiérrez, T., & Benetto, E. (2016). Framework and computational tool for the consideration of time dependency in Life Cycle Inventory: Proof of concept. *Journal of Cleaner Production*, *116*, 198–206. https://doi.org/10.1016/j.jclepro.2015.12.049

Tsoy, N., Steubing, B., Van Der Giesen, C., & Guinée, J. (2020). Upscaling methods used in ex ante life cycle assessment of emerging technologies: A review. *The International Journal of Life Cycle Assessment*, *25*(9), 1680–1692. https://doi.org/10.1007/s11367-020-01796-8

UNFCCC. (2025). *Nationally Determined Contributions (NDCs)*. https://unfccc.int/process-and-meetings/the-paris-agreement/nationally-determined-contributions-ndcs

Urbonavicius, M., Varnagiris, S., Mezulis, A., Lesnicenoks, P., Knoks, A., Richter, C., Milcius, D., Meirbekova, R., Gunnarsson, G., & Kleperis, J. (2024). Hydrogen from industrial aluminium scraps: Hydrolysis under various conditions, modelling of pH behaviour and analysis of reaction by-product. *International Journal of Hydrogen Energy*, *50*, 431–446. https://doi.org/10.1016/j.ijhydene.2023.09.065

US DOE. (2025). *Gaseous Hydrogen Delivery*. Energy.Gov. https://www.energy.gov/eere/fuelcells/gaseous-hydrogen-delivery

Van Beek, L., Hajer, M., Pelzer, P., Van Vuuren, D., & Cassen, C. (2020). Anticipating futures through models: The rise of Integrated Assessment Modelling in the climate science-policy interface since 1970. *Global Environmental Change*, *65*, 102191. https://doi.org/10.1016/j.gloenvcha.2020.102191

Van Der Giesen, C., Cucurachi, S., Guinée, J., Kramer, G. J., & Tukker, A. (2020). A critical view on the current application of LCA for new technologies and recommendations for improved practice. *Journal of Cleaner Production*, *259*, 120904. https://doi.org/10.1016/j.jclepro.2020.120904

Van Der Voet, E., Van Oers, L., Verboon, M., & Kuipers, K. (2019). Environmental Implications of Future Demand Scenarios for Metals: Methodology and Application to the Case of Seven Major Metals. *Journal of Industrial Ecology*, *23*(1), 141–155. https://doi.org/10.1111/jiec.12722

Van Vuuren, D. P., Edmonds, J., Kainuma, M., Riahi, K., Thomson, A., Hibbard, K., Hurtt, G. C., Kram, T., Krey, V., Lamarque, J.-F., Masui, T., Meinshausen, M., Nakicenovic, N., Smith, S. J., & Rose, S. K. (2011). The representative concentration pathways: An overview. *Climatic Change*, *109*(1–2), 5–31. https://doi.org/10.1007/s10584-011-0148-z

Van Vuuren, D. P., Kriegler, E., O'Neill, B. C., Ebi, K. L., Riahi, K., Carter, T. R., Edmonds, J., Hallegatte, S., Kram, T., Mathur, R., & Winkler, H. (2014). A new scenario framework for Climate Change Research: Scenario matrix architecture. *Climatic Change*, *122*(3), 373–386. https://doi.org/10.1007/s10584-013-0906-1

Vandepaer, L., Treyer, K., Mutel, C., Bauer, C., & Amor, B. (2019). The integration of long-term marginal electricity supply mixes in the ecoinvent consequential database version 3.4 and examination of modeling choices. *The International Journal of Life Cycle Assessment*, *24*(8), 1409–1428. https://doi.org/10.1007/s11367-018-1571-4

Vigon, B. W., & Jensen, A. A. (1995). Life cycle assessment: Data quality and databases practitioner survey. *Journal of Cleaner Production*, *3*(3), 135–141. https://doi.org/10.1016/0959-6526(94)00001-H

Walker, W. E., Harremoës, P., Rotmans, J., Van Der Sluijs, J. P., Van Asselt, M. B. A., Janssen, P., & Krayer Von Krauss, M. P. (2003). Defining Uncertainty: A Conceptual Basis for Uncertainty Management in Model-Based Decision Support. *Integrated Assessment*, *4*(1), 5–17. https://doi.org/10.1076/iaij.4.1.5.16466

Wan, B., Chen, W., Lu, T., Liu, F., Jiang, Z., & Mao, M. (2017). Review of solid state recycling of aluminum chips. *Resources, Conservation and Recycling*, *125*, 37–47. https://doi.org/10.1016/j.resconrec.2017.06.004

Weidema, B. (2003). Market Information in Life Cycle Assessment. In *Danish Environmental Protection Agency*.

Weidema, B. (2010). *Guidelines for application of deepened and broadened LCA Deliverable D 18 of work package 5 of the CALCAS project Authors*. https://www.semanticscholar.org/paper/Guidelines-for-application-of-deepened-and-LCA-D-18-Weidema/8be59252f6790328a6360d506df522de78bbce4c

Weidema, B., Bauer, C., Hischier, R., Mutel, C., Nemecek, T., Reinhard, J., Vadenbo, C. O., & Wernet, G. (2013). *Overview and methodology: Data quality guideline for the ecoinvent database version 3*.

Weidema, B. P., Pizzol, M., Schmidt, J., & Thoma, G. (2018). Attributional or consequential Life Cycle Assessment: A matter of social responsibility. *Journal of Cleaner Production*, *174*, 305–314. https://doi.org/10.1016/j.jclepro.2017.10.340

Weidema, Frees, N., & Nielsen, A.-M. (1999). Marginal production technologies for life cycle inventories. *The International Journal of Life Cycle Assessment*, *4*(1), 48–56. https://doi.org/10.1007/BF02979395

Weidema, & Wesnæs, M. (1996). Data quality management for life cycle inventories—An example of using data quality indicators. *Journal of Cleaner Production*, *4*(3–4), 167–174. https://doi.org/10.1016/S0959-6526(96)00043-1

Wender, B. A., Foley, R. W., Hottle, T. A., Sadowski, J., Prado-Lopez, V., Eisenberg, D. A., Laurin, L., & Seager, T. P. (2014). Anticipatory life-cycle assessment for responsible research and innovation. *Journal of Responsible Innovation*, *1*(2), 200–207. https://doi.org/10.1080/23299460.2014.920121

Wernet, G., Bauer, C., Steubing, B., Reinhard, J., Moreno-Ruiz, E., & Weidema, B. (2016). The ecoinvent database version 3 (part I): Overview and methodology. *The International Journal of Life Cycle Assessment*, *21*(9), 1218–1230. https://doi.org/10.1007/s11367-016-1087-8

Zamagni, A., Guinée, J., Heijungs, R., Masoni, P., & Raggi, A. (2012). Lights and shadows in consequential LCA. *The International Journal of Life Cycle Assessment*, *17*(7), 904–918. https://doi.org/10.1007/s11367-012-0423-x

Ziyadi, M., & Al-Qadi, I. L. (2019). Model uncertainty analysis using data analytics for life-cycle assessment (LCA) applications. *The International Journal of Life Cycle Assessment*, *24*(5), 945–959. https://doi.org/10.1007/s11367-018-1528-7

Zou, H., Chen, S., Zhao, Z., & Lin, W. (2013). Hydrogen production by hydrolysis of aluminum. *Journal of Alloys and Compounds*, *578*, 380–384. https://doi.org/10.1016/j.jallcom.2013.06.016

Appendix A — Stage D Substitution Governance Pack

A.1 Purpose and Scope

This appendix documents the governance rules used to convert recovered outputs from the modelled end-of-life (EoL) routes into Stage D substitution claims. Its purpose is to make explicit the minimum claim elements that sit behind every reported Stage D term, so that negative net results are interpreted as qualified avoided-burden claims rather than as intrinsic emissions attributes of the treatment routes themselves.

Appendix A records:

1. the governing interpretation of Stage D used in the thesis;
2. the route-by-route substitution-claim ledger;
3. the product-to-market mappings used for each creditable output;
4. the equivalence and conditioning gates that had to be met before credit was applied;
5. the hierarchy used when equivalence was partial or uncertain, including exclusion rules; and
6. the governance variants retained to test the robustness of the substitution claims.

This appendix should be read together with Appendix B, which documents the foreground route models and route-output quantities; Appendix C, which documents the frozen background databases; Appendix D, which documents market definitions and marginal supplier mixes; and Appendix F, which documents uncertainty and sensitivity analyses.

A.2 Role of Stage D

In this thesis, Stage D is treated as a conditional beyond-the-system-boundary consequence. A Stage D benefit is applied only when a recovered output from the foreground route is treated as capable of displacing a specific product or service in a defined receiving market. Accordingly, the Stage D term is governed by three linked components:

- the foreground route output, after all route-side losses and conditioning burdens;
- the equivalence rule, which determines whether the output qualifies for substitution and at what ratio; and
- the receiving-system definition, which specifies the displaced product, market scope, and marginal supplier representation.

Under this formulation, Stage D is not a property of the waste-treatment route alone. It is a property of the route plus the displaced market plus the rules used to connect them.

The thesis therefore reports decomposed values for C3–C4, Stage D, and net, and interprets net-negative results as qualified substitution claims whose credibility depends on the governance elements recorded below.

A.3 Governing Claim Elements

For a Stage D credit to be applied in this thesis, the following minimum claim elements had to be defined:

1. ***Displaced Product or Service:*** The specific good or service assumed to be displaced in the receiving system.
2. ***Market Boundary and Geography:*** The largest defensible receiving market that the recovered output could plausibly enter, defined by function, grade/specification where relevant, and consequential geography.
3. ***Marginal Supplier Representation:*** The supplier mix or avoided-production proxy used to value the displaced product or service in the relevant year/scenario.
4. ***Equivalence Gate:*** The rule determining whether the recovered output qualified for substitution, including required alignment in form, grade, service, purity, or product specification.
5. ***Conditioning Requirements:*** Any burdens and losses that had to be incurred before equivalence was deemed satisfied.
6. ***Substitution Ratio / Yield Mapping*****:** The conversion from route-output quantity to displaced-product quantity after all route-side losses, conditioning steps, and any quality-based realization logic.
7. ***Downscaling or Exclusion Rule***: The fallback rule used when full equivalence could not be defended.

These elements are recorded in the route-level claim ledger below (Table A-1). The ledger is organized by route and credited output. Where a claim element is constant across scenarios, it is stated once at the route level. Where it varies by year or scenario, the route entry points to the relevant discussion in Appendices B and D.

***Table A-1**: Stage-D substitution-claim ledger for retained routes*

Route / credited output	***Displaced product or service***	***Receiving market and geography***	***Marginal supplier representation***	***Equivalence gate***	***Conditioning required before credit***	***Substitution ratio / yield mapping***	***Downscaling / exclusion rule***
Landfill	None	None	None	No creditable output	None	None	No Stage D credit applied
Reuse	Aluminium non-structural façade-component replacement	Canada-wide façade-component/service market	Implemented through the foreground node AL_SD_credit_reuse_QC_ingot_plus_extrusion, combining avoided primary ingot and avoided extrusion burdens as a structured replacement proxy	The recovered component had to remain defensibly reusable in the same non-structural façade service class, with fit, geometry, and service duty compatible with the displaced application	Deconstruction, handling, cleaning, and degreasing only; no remelting or material transformation	Route basis = 1 kg aluminium. Reuse preparation burdens were scaled through the implemented degreasing conversion of 0.185185 m²/kg. The Stage D implementation then applied −1 kg canonical ingot credit plus −1 kg avoided extrusion per kg route basis	If direct reuse in the intended service class was not defensible, the reuse claim was excluded rather than retained at full credit
Conventional recycling	Wrought extrusion feedstock (billet-equivalent)	Canada-wide billet / extrusion-feedstock market	Contemporary implementation used the canonical Stage D primary-ingot proxy AL_credit_primary_ingot_IAI_NA_QC_elec; prospective interpretation follows the scenario-conditioned aluminium representation documented in Appendix D	Output had to remain within the defensible billet-equivalent / extrusion-feedstock class after route-side conditioning	Scrap preparation, refining/remelting, and any alloy-correction burdens embedded in the recycling route were incurred before Stage D	In the contemporary external_stageD implementation, the embedded credit was stripped and the recovered-yield scalar was inferred as 1.050300, with recycle_sub_ratio = 1.0, yielding 1.050300 kg avoided product per kg route basis	No additional discretionary downscaling was applied in the deterministic baseline beyond the implemented yield realization. If billet-equivalence were not defensible, the claim would be excluded
MS-FSC	Wrought extrusion feedstock (billet-equivalent)	Canada-wide billet / extrusion-feedstock market	Route-specific avoided primary-aluminium credit nodes, including baseline and inert-overlay forms, as documented in Appendices B and D	Output had to pass the billet-equivalence and quality/eligibility gate used to justify displacement of extrusion-grade billet/feedstock	Shredding, degreasing, consolidation, trimming, and route-side realization were incurred before Stage D	Deterministic baseline realization was governed by the implemented route yields: shred yield × FSC retained yield × substitution ratio = 0.80 × 0.952 × 1.0 = 0.7616 kg credited billet-equivalent output per kg qualifying chain input. Any pass-share reduction was treated as a robustness/sensitivity variant rather than the deterministic default	Non-qualifying output receives no Stage D credit. Partial pass shares were treated as reduced-claim variants
Hydrolysis — H_2	Gaseous hydrogen represented through a low-pressure	Canada-wide hydrogen receiving system	Contemporary: foreground-constructed hydrogen central case documented in Appendix D. Prospective: scenario-conditioned hydrogen representation implemented through the h2_market_low_pressure_proxy_C	Hydrogen had to be purified to the product basis represented by the low-pressure hydrogen receiving-system proxy	Hydrolysis reaction and PSA purification were included before credit. No separate medium-pressure compression basis was	Credited hydrogen was calculated from reacted aluminium and PSA recovery: $m_{H_{2,credited}} = m_{Al,reacted} \times \frac{3M_{H_2}}{2M_{Al}} \times R_{PSA}$ Using the deterministic baselines, this corresponds to approximately 0.0587 kg H_2/kg	If the hydrogen product basis could not be defended, the claim was excluded

	hydrogen market proxy		A_prospect_locpref__<scenario> proxy family		retained in the final implementation	gate scrap in the 2025 case and 0.0860 kg H_2/kg gate scrap in the 2050 prospective case	
Hydrolysis — $Al(OH)_3$	Aluminium hydroxide supply	Global background market proxy	Implemented through the aloh3_market_proxy foreground proxy family, linked to a global aluminium-hydroxide market basis	The recovered solids had to remain compatible with the background aluminium-hydroxide supply proxy	No additional drying or separate conditioning step was imposed in the final implementation. Water associated with the reaction was assumed to evaporate or separate to the biosphere rather than requiring a separate conditioning burden on the $Al(OH)_3$ product	Credited quantity was based on stoichiometric recovery from reacted aluminium: $$m_{Al(OH)_3} = m_{Al,reacted} \times \frac{M_{Al(OH)_3}}{M_{Al}}$$ This gives approximately 1.97 kg $Al(OH)_3$/kg gate scrap for the 2025 contemporary baseline and 2.33 kg $Al(OH)_3$/kg gate scrap for the 2050 prospective baseline	If the solids stream could not be defended as market-compatible, the claim was excluded

A.3.1 Interpretation Note

Table A.1 is the thesis's minimum claim register. It does not reproduce every mass-balance step or route parameter; those remain in Appendix B. Instead, it records the rule by which route output becomes an avoided-burden claim in Stage D.

A.4 Product-to-Market Mapping

Table A-1 is the thesis's minimum claim register. It does not reproduce every mass-balance step or route parameter; those remain in Appendix B. Instead, it records the rule by which route output becomes an avoided-burden claim in Stage D.

Table A-2: *Product-to-market mapping for creditable outputs*

Route	***Foreground output***	***Output basis in foreground model***	***Mapped receiving product***	***Receiving market basis***	***Notes***
Reuse	Reusable aluminium component output	1 kg aluminium route basis with preparation burdens scaled through the degreasing conversion	Aluminium non-structural façade-component replacement	Canada-wide	The market boundary remained a façade-component/service claim, but the final implementation used a combined ingot-plus-extrusion proxy because a directly observable reuse market time series was not available
Conventional recycling	Refined/remelted metallic output	1 kg aluminium route basis after route-side realization	Wrought extrusion feedstock (billet-equivalent)	Canada-wide	Implemented through an avoided primary-ingot proxy in the contemporary case
MS-FSC	Consolidated billet-equivalent output	Route-specific chain input basis, realized after shred and retained-yield losses	Wrought extrusion feedstock (billet-equivalent)	Canada-wide	Quality-gated billet-equivalent substitution
Hydrolysis	Purified gaseous hydrogen	Route output after hydrolysis realization and PSA recovery	Gaseous hydrogen represented through a low-pressure market proxy	Canada-wide	Final implementation used a low-pressure hydrogen proxy basis
Hydrolysis	Recovered aluminium hydroxide solids	Recovered solids after stoichiometric realization	Aluminium hydroxide supply	Global background proxy	No additional conditioning step was imposed in the final implementation
Landfill	None	None	None	None	No mapped product

A.4.1 Coproduct Handling

Only hydrolysis generates multiple distinct candidate credited outputs in the final thesis model. These outputs were not collapsed into a single generic credit. Instead:

- hydrogen was mapped to the hydrogen receiving-system proxy; and
- aluminium hydroxide was mapped separately to the global aluminium-hydroxide market proxy.

This separation is important because the two outputs differ in product form, market structure, and receiving-system carbon intensity. Their Stage D contributions therefore remain independently governed even when reported together at route level.

A.5 Implementation Note on Equivalence, Conditioning and Claim Failure

Detailed route-specific equivalence and conditioning logic is described in the main text and in Appendix B. In the final thesis model, that logic was implemented directly through the foreground route construction and Stage D mapping rather than through a separate post hoc downscaling hierarchy.

Accordingly, no general downscaling ladder was retained as an active deterministic feature of the final model. Routes either met the implemented claim basis and received the corresponding Stage D credit, or they did not. Landfill therefore received no Stage D credit by definition. For MS-FSC, only qualifying billet-equivalent output was credited, while reduced pass-share realization was examined separately as a robustness variant. For hydrolysis, hydrogen and aluminium hydroxide were credited separately against their respective receiving-system proxies only where the implemented product basis was satisfied. For reuse and conventional recycling, the Stage D claim followed the fixed route definitions and receiving-system proxies described above.

A.6 Governance Variance Retained in Robustness Testing

Because many net results in the thesis are Stage D-dominant, robustness is not established by deterministic route comparison alone. The thesis therefore retained a bounded set of governance variants to test whether conclusions depended materially on receiving-system definition and claim realization. The retained variants are documented below in Table A-3.

Table A-3*: Governance-Variant Register*

Variant family	***Applied to***	***Variant description***	***Location***
Import-enabled supplier variant	Primary aluminium displacement	Tests whether screened import participation materially alters the contemporary primary aluminium marginal mix	Appendix D

CRR-screened / alternative supplier-screen variant	Primary aluminium displacement	Tests bounded changes in marginal-supplier eligibility under alternative screening logic	Appendix D
Prospective aluminium technology overlay	Primary aluminium displacement in 2050	Retains contemporary supplier identity while varying production mechanism between conventional and inert-anode / ELYSIS-type supply	Appendix D
Qualification / pass-share realization	MS-FSC	Tests the effect of less-than-full qualifying output on Stage D realization	Appendix F
Reduced route realization	Hydrolysis	Tests how route-performance terms affect the strength of the Stage D hydrogen claim under fixed receiving-system assumptions	Appendix F

A.7 Interpretation Note

The variants recorded here do not imply that every claim element was redefined in every robustness test. Rather, they identify the limited set of changes that materially affect the strength, scope, or realization of the substitution claim and therefore matter for interpretation of Stage D-dominant net results.

Appendix B — Foreground Models, Proxies, Parameters, and Regionalization

B.1 Purpose & Scope

This appendix documents the foreground model layer used to implement the retained end-of-life (EoL) routes assessed in the thesis. In addition to the implemented route models, it records the targeted-review longlist and screening register used to identify candidate emerging aluminium EoL technologies prior to final route selection. It records the representative curtainwall material schedule, the aluminium-route accounting basis used in the foreground wrappers, the route-wrapper logic for landfill, reuse, conventional recycling, multi-step friction stir consolidation (MS-FSC), and aluminium hydrolysis, the proxy strategy used where direct database representations were unavailable, the principal deterministic parameters used to instantiate the baseline models, the main utility-regionalization choices applied to foreground processes, and the treatment of non-priority materials.

The appendix is intentionally selective. It reports the information required to understand, reproduce, and audit the foreground model structure used in the thesis, but it does not reproduce every activity-build script, exchange-edit log, or full parameter manifest. Those implementation-level materials remain archived in the accompanying repository.

This appendix should be read together with Appendix A, which documents the Stage D substitution-claim governance; Appendix D, which documents market boundaries and marginal supplier mixes; and Appendix F, which documents uncertainty propagation and sensitivity analyses.

B.2 Targeted Review and Screening Register for Emerging Aluminium EoL Technologies

Before the foreground route set was finalized, a targeted review of the uploaded literature was used to identify a broader longlist of candidate emerging aluminium end-of-life technologies. This review served a screening function rather than an implementation function: its purpose was to identify which routes were sufficiently mature, sufficiently described, and sufficiently compatible with consequential substitution logic to justify foreground modelling in the thesis.

The longlist was developed from the uploaded route-specific and review literature, including the solid-state recycling review material and the individual studies on friction-stir consolidation and aluminium hydrolysis. Candidate routes were screened against the admission policy defined in Chapter 3: (i) traceable LCI-compatible parameterization, (ii) plausible displaced-product intent, and (iii) suitability for auditable Stage D implementation. Routes that helped define the broader emerging-technology landscape but did not meet the final implementation threshold were retained here as contextual candidates rather than discarded from the record.

Table B-1 summarizes the resulting longlist and indicates how each candidate was treated in the thesis.

***Table B-1**: Targeted-review longlist and screening register for candidate emerging aluminium EoL technologies*

Candidate route / technology family	***Representative basis in reviewed/uploaded PDFs***	***Indicative output / displaced-product intent***	***Indicative maturity for thesis framing***	***Screening outcome in thesis***
Direct chip-to-profile hot extrusion	(D. Cooper, 2013; Duflou et al., 2015)	Extruded aluminium profile or billet-equivalent semi-finished product	Mid-TRL	Valid longlist candidate; not retained in final case-study route set
ECAP-assisted hot extrusion of aluminium chips	(Duflou et al., 2015)	Billet/profile-grade solid-state recycled product	Mid-TRL	Valid longlist candidate; not retained as a standalone route
Screw extrusion of aluminium scrap/chips	(Duflou et al., 2015)	Continuous solid-state extrudate / profile-type product	Mid-TRL	Valid longlist candidate; not retained as a standalone route
Spark plasma sintering (SPS) of aluminium scrap/chips	(Duflou et al., 2015)	Dense near-net-shape consolidated aluminium product	Mid-TRL	Valid longlist candidate, but not retained because the displaced-product basis differs from the billet-oriented routes used in the thesis
Single-step friction stir consolidation (FSC)	(Ingarao et al., 2023)	Billet-equivalent consolidated aluminium from chips	Mid-TRL	Valid longlist candidate; superseded in final selection by multi-step FSC
Multi-step friction stir consolidation (MS-FSC)	(Ingarao et al., 2023; Demirci et al., 2025)	Billet-equivalent consolidated aluminium	Mid-TRL	Retained in final thesis route set

		suitable for extrusion-feedstock substitution		
Functionally graded / multi-material FSC billets	(Latif, Ingarao, Gucciardi, et al., 2022)	Upcycled graded billet rather than conventional billet-equivalent feedstock	Low-to-mid TRL	Retained as contextual evidence for FSC flexibility, but not as a standalone displaced-product route
Hydrolysis of aluminium scrap for hydrogen generation	(Urbonavicius et al., 2024)	Merchant gaseous hydrogen with aluminium hydroxide co-product	Low TRL	Retained in final thesis route set
Hydrolysis using alkaline promoters / pretreatments	(Petrovic & Thomas, 2011; Zou et al., 2013)	Same hydrogen-service intent with altered reaction performance	Low TRL	Used as supporting evidence for hydrolysis parameterization; not retained as a separate route
Hydrolysis using alternative aqueous streams / waste streams	(Deebika & Saravanakumar, 2023)	Hydrogen generation from waste aluminium under modified reaction medium	Low TRL	Contextual longlist candidate; not retained because system expansion and transferability become substantially less direct
Improved remelting or pre-melting consolidation variants for chip recovery	(D. R. Cooper et al., 2018)	Improved recycling yield rather than a distinct non-melting route	Adjacent process-improvement class	Used as contextual evidence for conventional-recycling yield logic; not retained as a separate emerging route

B.3 Representative Assembly and Aluminium Subsystem Basis

B.3.1 Representative Curtainwall System

The representative curtainwall panel was based on the RAICO THERM+ H-I timber façade with a 76 mm timber mullion/transom width ((RAICO, 2026). Table B-1 reports the full material schedule for a typical office building mass timber curtainwall façade,

derived from the DOE prototype medium office building (DOE, 2023), as well as the normalization to the representative 1 m² assembly basis from the full façade surface area.

Table B-2: *Full bill of materials for the representative curtainwall panel (~33m² surface area)*

Component	***Material Name***	***Qty***	***Units***	***Quantity (kg)***	***Density (kg/m³)***	***Quantity (m³)***	***Normalized Quantity (kg/m²)***
Transom & Mullion	Cross-Laminated Timber (CLT)	0.58465765	m3	280.635672	480	0.58465765	8.504111273
Snap Cover	CLT	0.099980003	m3	47.99040156	480	0.099980003	1.454254593
Pressure Plate	Aluminium	4.633092363	m2	38.91797585	2800	0.013899277	1.179332602
Exterior Gasket	Ethylene Propylene Diene Monomer (EPDM)	7.153873983	kg	7.153873983	900	0.007948749	0.21678406
Insulating Block	Polyamide	41.26865956	kg	41.26865956	1100	0.037516963	1.250565441
Mullion/Transom Gasket	EPDM	25.32193989	kg	25.32193989	900	0.028135489	0.767331512
Thermal Base Profile	Polyamide	15.21468223	kg	15.21468223	1100	0.013831529	0.461050977
Spacer	Aluminium	7.49240879	kg	7.49240879	2800	0.00267586	0.227042691
Glass	Double Glazed Flat/Float Glass	21.23300416	m2	523.208558	2500	0.209283423	15.85480479
Coating	Metallic Oxide	N/A	N/A	–	–	–	–

Desiccant	Silica Gel	3.70656202	kg	3.70656202	2200	0.001684801	0.112320061
Sealant	Silicone	2.812263107	kg	2.812263107	2329	0.001207498	0.085220094
Gas Infill	Argon	0.79330807	kg*	0.569250911	1.7	0.334853477	0.017250028
High Density XPS Block	Extruded Polystyrene Insulation (XPS)	0.270144246	kg	0.270144246	24.8	0.010892913	0.008186189
Insulation	XPS	4.36718953	kg	4.36718953	24.8	0.176096352	0.132339077
Backpan Insulation	Fibre Glass	5.559644342	kg	5.559644342	12	0.463303695	0.168474071
Front Face	Aluminium	71.78602054	kg	71.78602054	2800	0.025637864	2.175333956
Membrane	Polyethylene	1.900558048	kg	1.900558048	950	0.002000587	0.057592668
Blocking	Softwood	0.050203825	m3	25.10191248	500	0.050203825	0.760664015
Plywood	Softwood	0.095832316	m3	47.916158	500	0.095832316	1.452004788
Framing	Softwood	0.02368477	m3	11.8423848	500	0.02368477	0.358860145
Backpan/Interior Finish	Gypsum	67.87690672	kg	67.87690672	700	0.09696701	2.056875961
Vertical Flashing	Aluminium	1.021513595	kg	1.021513595	2800	0.000364826	0.030954957
Horizontal Flashing	Aluminium	1.747371172	kg	1.747371172	2800	0.000624061	0.052950642
Setting Block	Silicone	9.280807245	kg	9.280807245	2800	0.003314574	0.281236583
Glass Carrier	Fibre Glass	0.025452519	kg	0.025452519	12	0.002121043	0.000771288

Table B-1 should be interpreted as the full assembly mass schedule used to anchor the representative façade basis. The directly listed aluminium-bearing components in Table B-1 are the pressure plate, spacer, front face, and vertical and horizontal flashings. Taking their associated quantities (kg/m^2) yields:

Equation B-1

$$1.179 + 0.2270 + 2.175 + 0.0310 + 0.0530 \approx 3.67\ kg/m^2$$

This route-model aluminium basis is the scalar used in all base-route, MS-FSC, and hydrolysis LCIA runs.

B.4 Route Wrappers and Reference Flow Definitions

B.4.1 General Wrapper Logic

The foreground model was organized around five retained route wrappers:

1. landfill,
2. reuse,
3. conventional recycling,
4. MS-FSC,
5. aluminium hydrolysis.

Each wrapper separates the burden-side C3–C4 process chain from the output conditioning and Stage D realization logic. Final results are reported on the 1 m² functional unit using the aluminium subsystem basis of 3.67 kg Al/m².

Two reference-flow conventions were used in the executed foreground models:

1. A 1 kg aluminium-service basis for the base aluminium routes (landfill, reuse, conventional recycling), later scaled to 3.67 kg Al; and
2. A gate-basis convention for the emerging routes, where one unit demanded of the route wrapper represents 1 kg aluminium scrap treated at the route gate.

This distinction matters for auditing the runners. In the hydrolysis models, the C3–C4 activity and the Stage D offset activities are demanded on the same gate basis, so the joint 3.67-unit demand ensures that route burdens and route credits correspond to the same treated input quantity. The MS-FSC chain is similarly gate-aligned, although Stage D realization is further governed by the route's billet-equivalence logic and any quality or pass-share assumptions applied downstream.

B.4.2 Route Summary

Table B-3: *Foreground route wrapper summary*

Route	***Representative wrapper / node logic***	***Route-entry basis***	***Main C3–C4 operations***	***Creditable output(s)***	***Stage D gate***
Landfill	AL_RW_landfill_C3C4_CA / AL_RW_landfill_NET_CA	1 kg Al	Collection, handling, disposal	None	No Stage D
Reuse	AL_RW_reuse_C3_CA / AL_RW_reuse_NET_CA	1 kg Al	Preparation and cleaning via degreasing proxy	Like-for-like façade-component service	After reuse eligibility is met
Conventional recycling	AL_RW_recycling_postcons_refiner_C3C4_CA / AL_RW_recycling_postcons_NET_CA	1 kg Al	Scrap preparation, refiner/remelting treatment, alloy correction	Billet-equivalent feedstock	After billet-equivalence is met
MS-FSC	Gate diversion → shredding → degreasing → consolidation → Stage D credit wrapper	1 kg chain-gate scrap / prepared scrap basis	Shredding, degreasing, solid-state consolidation, trimming	Billet-equivalent output	After pass / quality gate is met
Hydrolysis	Gate scrap → preparation → hydrolysis route → Stage D H_2 / $Al(OH)_3$ offsets	1 kg gate scrap basis	Preparation, reaction, liquor handling, PSA purification, solids recovery	Hydrogen and $Al(OH)_3$	After product-specific conditioning is met

B.5 Proxy Modelling Strategy

B.5.1 Minimum-Change Proxy Rule

Where no direct database activity existed for a required unit operation, or where the available activity did not align with the intended function, basis, or geography, a proxy was created using a minimum-change rule. Under this rule, the closest suitable background activity was cloned and only the exchanges required to represent the intended unit operation were modified. The resulting foreground activities therefore preserve the upstream realism of the source activity while keeping the foreground edits bounded and auditable.

Across the developed scripts, provided in the repository, five recurring proxy operations were used:

- stripping embedded consequential credits from copied treatment activities;

- localizing electricity and selected utilities;
- rescaling copied activities to the intended process basis;
- constructing explicit foreground-only processing nodes where no suitable ecoinvent process existed; and
- isolating receiving-system credit proxies so that avoided burdens were transparent and separately auditable.

Table B-4 provides the detailed proxy registry.

***Table B-4**: Proxy registry*

Proxy group	***Representative foreground nodes***	***Typical source activity family***	***Main edits***	***Used in***
Scrap-gate and burdens-only treatment proxies	al_scrap_postconsumer_CA_gate, al_scrap_postconsumer_CA_gate_FSC, AL_UP_refiner_postcons_NO_CREDIT_CA	Post-consumer aluminium scrap treatment / refiner treatment	Copied treatment burdens retained; negative technosphere exchanges stripped where embedded consequential credits had to be isolated; route basis renormalized	Conventional recycling, MS-FSC, hydrolysis
Preparation proxies	AL_UP_degreasing_CA, FSC_degreasing_CA, al_scrap_shredding_for_hydrolysis_CA	Metal-part degreasing and scrap preparation activities	Electricity scaling, yield scaling, and where needed surface-area conversion or shred/prep normalization	Reuse, MS-FSC, hydrolysis
Hydrolysis service proxies	di_water_CA, wastewater_treatment_unpolluted_CAe, naoh_CA_proxy, naoh_electrolyte_solution_CA, h2_purification_psa_service_CA	Water production, wastewater treatment, sodium hydroxide market, biogas purification by PSA	Utility relinking, quantity scaling, electrolyte make-up construction, and gate-basis normalization	Hydrolysis
MS-FSC process	FSC_consolidation_CA and scenario equivalents	Foreground-added process node	Explicit electricity, lubricant, shred-yield,	MS-FSC

proxy / explicit foreground node		informed by measured processing inputs; supported by preparation-chain proxies	and retained-yield terms; no reliance on an off-the-shelf consolidation activity	
Receiving-system and credit proxies	AL_credit_primary_ingot_IAI_NA_QC_elec, AL_UP_avoided_impact_extrusion_CA, hydrogen LP market proxies, Al(OH)(_3) market proxies, SMR provider clone	Aluminium ingot market / primary ingot production, impact extrusion, hydrogen market, aluminium hydroxide market, hydrogen production provider	Localization, supplier forcing, electricity relinking, and/or credit isolation to create explicit avoided-burden nodes	Reuse, conventional recycling, MS-FSC, hydrolysis
Non-priority material treatment proxies	Per-material C3–C4 and Stage D wrappers for CLT, SPF, glass, XPS, polyethylene, polyamide, silicone, silica gel, EPDM, gypsum, and fibre glass	Waste wood incineration, waste glass landfill, waste polystyrene landfill, waste polyethylene landfill, mixed plastic landfill, MSW landfill, waste rubber incineration, waste gypsum landfill, waste mineral wool recycling	Template copy, split of positive versus negative technosphere exchanges, electricity/utility localization, scenario-specific wrapper generation where required	Non-priority material layer

B.6 Key Foreground Parameters

B.6.1 Parameter Reporting Principle

Only the principal deterministic parameters needed to interpret the baseline foreground models are reported here. Full script variable names, uncertainty ranges, and Monte Carlo distribution families remain documented in the repository and, where relevant to uncertainty propagation, in Appendix F.

The most important distinction is between mature base routes and emerging routes. The mature base routes rely primarily on wrapper structure and explicit separation of C3–C4 versus Stage D logic. The emerging routes rely on a compact set of foreground parameters that control energy demand, yields, and the quantity of creditable output. The baseline parameters for MS-FSC and hydrolysis are reported in Table B-5 and Table B-6.

B.6.2 MS-FSC Baseline Parameters

Table B-5*: Principal MS-FSC baseline parameters*

Parameter	***2025 contemporary***	***2050 prospective***	***Interpretation***
Shredding electricity	0.30 kWh/kg scrap	0.30 kWh/kg scrap	Burden-side shredding electricity
Shred yield	0.80	0.80	Usable shredded feedstock from prepared scrap
FSC retained yield	0.952	0.952	Billet output per kg shredded/degreased feed
Consolidation energy (A)	0.267 MJ per 20 g billet	0.267 MJ per 20 g billet	Productive consolidation energy
Transition energy (B)	0.355 MJ per 20 g billet	0.355 MJ per 20 g billet	Transition / ramp overhead
Include transition term?	True	False	Determines whether (B) is included in baseline electricity
Lubricating oil	0.02 kg/kg billet	0.02 kg/kg billet	Process consumable
Stage D substitution ratio	1.0	1.0	kg avoided primary ingot per kg billet-equivalent output

MS-FSC was treated as a mid-TRL route whose core process performance is already comparatively constrained relative to the lower-TRL hydrolysis pathway. As shown in Table B-4, the route-specific parameters most directly tied to the emerging technology itself, FSC retained

yield, productive consolidation energy, transition energy, and lubricating-oil demand, were either held constant across time horizons or varied only through the explicit inclusion or exclusion of the transition term in the baseline electricity formulation. By contrast, shredding electricity and shred yield were treated as upstream preparation parameters associated with more established scrap-handling operations rather than with the novel consolidation step itself. Accordingly, they were not interpreted as primary technology-improvement levers for MS-FSC in this study and were therefore not prioritized for prospective improvement or uncertainty sampling. This parameterization reflects the view that the main prospective change in MS-FSC lies not in radical changes to billet yield or consumable demand, but in the reduction of laboratory-specific overheads, particularly the transition or ramp-energy burden, as the process moves toward more stable industrial operation.

B.6.3 Hydrolysis Baseline Parameters

Table B-6*: Principal hydrolysis baseline parameters*

Parameter	***2025 contemporary***	***2050 prospective***	***Interpretation***
(f_Al) / metallic Al fraction	1.00	1.00	Metallic Al available in prepared feed
(X_Al) / reacted fraction	0.85	0.95	Effective Al conversion
Liquor intensity	250 L/kg Al	150 L/kg prepared scrap	Working liquor burden basis
NaOH molarity	0.240 M	0.240 M	Working chemistry concentration
Liquor make-up / solvent loss	1.00 make-up fraction	0.20 solvent-loss fraction	Contemporary and prospective scripts parameterize make-up differently
Preparation yield (Y_prep)	0.80	0.85	Prepared feed per kg gate scrap or inverse gate requirement
Preparation electricity	0.0504 kWh/kg prepared	0.0504 kWh/kg prepared	Scrap preparation electricity
PSA recovery (R_PSA)	0.77	0.95	Usable H_2 after purification
Stoichiometric water source	liquor pool	liquor pool	Water accounting convention

The hydrolysis baseline parameters distinguish a contemporary, laboratory-style implementation from a more engineered prospective configuration while preserving the same underlying process topology. The 2025 case represents a comparatively dilute, loss-intensive, and less controlled operating state, reflected in lower effective Al conversion, higher liquor intensity, full liquor make-up, lower preparation yield (aligned with the values reported by Urbonavicius et al., 2024), and lower PSA recovery. The 2050 case does not assume a fundamentally different chemistry; rather, it represents improved process control and circulation efficiency within the same hydrolysis framework. This is reflected in the higher reacted fraction, reduced working-liquor burden, partial solvent-loss treatment instead of full make-up, modestly improved preparation yield, and substantially higher hydrogen purification recovery. NaOH molarity and the stoichiometric water-accounting convention are held constant across both cases to avoid embedding unverified chemistry changes into the prospective baseline. Accordingly, the prospective parameter shift should be interpreted as a bounded engineering maturation of the route rather than as a speculative redesign of the process itself.

B.6.4 Base Routes

The mature base routes rely less on an extensive foreground parameter set and more on wrapper structure and embedded route logic. Their major foreground scalars, such as preparation burdens or degreasing conversions, are documented in the repository.

B.7 Utility Regionalization and Provider Alignment

Utility regionalization was applied selectively where electricity or other utilities materially influenced route burdens and where a geography-consistent provider could be justified. Across the reviewed scripts, the following utility reference products were repeatedly targeted for localization or provider replacement:

- tap water,
- wastewater, average,
- heat, district or industrial, natural gas,
- heat, district or industrial, other than natural gas,
- light fuel oil,
- heavy fuel oil,
- lubricating oil.

Canada was generally preferred first, followed by province-specific or Canada-adjacent matches where relevant to the route logic, then North American or broader fallback providers such as Rest of World (RoW) or Global (GLO). However, the exact preference order differed slightly by route family. Table B-7 details the provider and utilities general to each route.

B.7.1 Summary of Route-Specific Provider Alignment

Table B-7: *Utility and provider-alignment summary*

Route family	***Burden-side electricity***	***Credit-side provider basis***	***Other utility relinking***	***Notes***
Base aluminium routes	Contemporary or scenario-consistent background electricity	Receiving-system-specific avoided-product proxy	Limited	Base routes rely mainly on wrapper structure rather than extensive utility relinking
MS-FSC	Localized burden-side electricity	Avoided-ingot proxy with route-appropriate provider basis	Limited	Burden and credit electricity were intentionally separated
Hydrolysis	Localized foreground electricity	Product-specific H_2 and $Al(OH)_3$ receiving systems	Water, wastewater, NaOH, PSA-related services	Hydrolysis required the most explicit utility alignment
Non-priority materials	Simplified localized wrapper electricity	Embedded credit logic where applicable	Minimal	Implemented as simplified wrappers rather than bespoke route chains

B.8 Worked Examples of Route-Level Scaling

The purpose of these worked examples is to illustrate how the foreground route parameters translate a unit quantity of route-entry material into the intermediate or creditable outputs used later in the burden and Stage D calculations. These examples are not intended to reproduce the full model logic from the scripts, but rather to show the main mass-balance relationships that govern each route at the unit-process level.

B.8.1 Example 1 — MS-FSC Billet Realization from Prepared Scrap

For the contemporary MS-FSC baseline, the billet-equivalent output is governed primarily by the shred yield and the retained yield of the consolidation step. Using the central parameters from Table B.4:

- shred yield = 0.80
- FSC retained yield = 0.952

For 1 kg of prepared scrap entering the MS-FSC chain:

Equation B-2

$$m_{\text{shredded}} = 1.0 \times 0.80 = 0.80 \text{ kg}$$
$$m_{\text{billet}} = 0.80 \times 0.952 = 0.7616 \text{ kg billet-equivalent}$$

Thus, the contemporary MS-FSC chain realizes:

$$0.7616 \text{ kg billet-equivalent per kg prepared scrap}$$

This example illustrates the main physical logic of the route: only a fraction of the prepared scrap remains as usable shredded feed, and only a fraction of that feed is retained as billet-equivalent output after consolidation and trimming. In the model, this billet-equivalent mass is the quantity that proceeds to the Stage D equivalence gate for substitution against the displaced billet market.

B.8.2 Example 2 — MS-FSC Electricity Formulation

The electricity burden for MS-FSC is composed of the burden from shredding plus the burden from friction-stir consolidation itself. The literature-derived consolidation term is split into:

- productive consolidation energy, $A = 0.267$ MJ per 20 g billet
- transition or ramp overhead, $B = 0.355$ MJ per 20 g billet

Converted to a kilogram basis:

Equation B-3

$$A = \frac{0.267}{0.020} = 13.35 \text{ MJ/kg billet}$$

Equation B-4

$$B = \frac{0.355}{0.020} = 17.75 \textit{ MJ/kg billet}$$

Using:

$$1 \text{ kWh} = 3.6 \text{ MJ}$$

gives:

Equation B-5

$$A = \frac{13.35}{3.6} = 3.708 \text{ kWh/kg billet}$$

Equation B-6

$$B = \frac{17.75}{3.6} = 4.931 \textit{ kWh/kg billet}$$

Therefore, the contemporary baseline includes both terms:

Equation B-7

$$E_{\mathrm{FSC},2025} = A + B = 3.708 + 4.931 = 8.639 \text{ kWh/kg billet}$$

By contrast, the prospective baseline excludes the transition term for the deterministic analysis:

Equation B-8

$$E_{\mathrm{FSC},2050} = A = 3.708 \text{ kWh/kg billet}$$

This formulation clarifies that the principal prospective improvement represented for MS-FSC is not a change in the productive consolidation requirement itself, but the removal of the laboratory-specific transition overhead from the baseline electricity burden.

B.8.3 Example 3 — Hydrolysis Reacted Aluminium from Prepared Feed (Contemporary)

For the contemporary hydrolysis baseline, the chain converts gate-entry scrap to prepared feed and then to reacted metallic aluminium according to the central parameters:

- preparation yield, $Y_{\mathrm{prep}} = 0.80$
- metallic Al fraction, $f_{\mathrm{Al}} = 1.00$
- reacted fraction, $X_{\mathrm{Al}} = 0.85$

For 1 kg of gate-entry scrap:

Equation B-9

$$m_{\mathrm{prepared}} = 1.0 \times Y_{\mathrm{prep}} = 1.0 \times 0.80 = 0.80 \text{ kg prepared feed}$$
$$m_{\mathrm{Al,available}} = 0.80 \times f_{\mathrm{Al}} = 0.80 \times 1.00 = 0.80 \text{ kg metallic Al}$$
$$m_{\mathrm{Al,reacted}} = 0.80 \times X_{\mathrm{Al}} = 0.80 \times 0.85 = 0.68 \text{ kg reacted Al}$$

Thus, the contemporary hydrolysis model realizes 0.68kg of reacted aluminium per kg of gate-entry scrap. This reacted aluminium mass is the key process-basis quantity from which downstream hydrogen and aluminium hydroxide outputs are calculated in the foreground model.

B.8.4 Example 4 — Hydrolysis Liquor and Chemistry Burden Basis (Contemporary)

The hydrolysis route also carries a substantial working-liquor burden, represented differently in the contemporary and prospective baselines.

For the 2025 contemporary case, the liquor intensity is 250 L/kg of aluminium, in alignment with Urbonavicius et al. (2024). Using the reacted aluminium quantity from Equation B-9:

Equation B-10

$$V_{\mathrm{liquor},2025} = 0.68 \times 250 = 170 \text{ L liquor per kg gate-entry scrap}$$

For the 2050 prospective case, the route is represented as a more controlled and less dilute configuration. The central values are:

- preparation yield, $Y_{\mathrm{prep}} = 0.85$
- metallic Al fraction, $f_{\mathrm{Al}} = 1.00$
- reacted fraction, $X_{\mathrm{Al}} = 0.95$
- liquor intensity = 150 L/kg prepared scrap

For 1 kg of gate-entry scrap:

Equation B-11

$$m_{\mathrm{prepared,2050}} = 1.0 \times 0.85 = 0.85 \text{ kg prepared feed}$$
$$m_{\mathrm{Al,reacted,2050}} = 0.85 \times 1.00 \times 0.95 = 0.8075 \text{ kg reacted Al}$$

The working-liquor requirement is then:

Equation B-12

$$V_{\mathrm{liquor,2050}} = 0.85 \times 150 = 127.5 \text{ L liquor per kg gate-entry scrap}$$

This comparison shows that the prospective hydrolysis baseline relies upon the assumption of improvement of both conversion and liquor efficiency. The route therefore treats more aluminium while carrying a lower working-liquor burden per kilogram of gate-entry material.

B.9 Non-Priority Material Results and Interpretation

The non-priority materials were not omitted from the assessment. They were implemented as simplified localized foreground wrappers and evaluated under the same C3–C4 / Stage D split used elsewhere in the thesis. Their role in the appendix is not to establish a second focal result stream, but to confirm that the remainder of the façade material schedule was carried through the accounting consistently and that its contribution is modest relative to the aluminium subsystem. Table B-8 reports the treatment assumptions used for the non-priority materials while Table B-9 summarizes the resulting net GWP100 outcomes across all scenarios.

Table B-8: *Non-priority material treatment*

Material family	***Proxy treatment intent***	***Typical template family***
CLT (mass timber)	Energy-recovery proxy	Waste wood, untreated, municipal incineration
SPF softwood	Energy-recovery proxy	Waste wood, untreated, municipal incineration
Glass	Landfill proxy	Waste glass, sanitary landfill
XPS	Landfill proxy	Waste polystyrene, sanitary landfill
Polyethylene	Landfill proxy	Waste polyethylene, sanitary landfill
Polyamide	Mixed-plastic / MSW landfill proxy	Waste plastic mixture, sanitary landfill
Silicone	MSW landfill proxy	Municipal solid waste, sanitary landfill
Silica gel	MSW landfill proxy	Municipal solid waste, sanitary landfill
EPDM	Incineration proxy	Waste rubber, unspecified, municipal incineration
Gypsum	Landfill proxy	Waste gypsum, sanitary landfill
Fibre glass	Recycling proxy	Waste mineral wool, recycling

Table B-9: *Selected non-priority material net GWP100 results (kg CO_2e/m^2)*

Material family	***Contemporary***	***SSP1-VLLO***	***SSP2-M***	***SSP5-H***
CLT	−5.81	−5.96	−5.99	−4.46
SPF softwood	−1.50	−1.54	−1.55	−1.15
EPDM	+2.25	+2.21	+2.21	+2.06
Polyamide	+0.18	+0.18	+0.18	+0.18
Glass	+0.18	+0.15	+0.16	+0.16
Silicone	+0.24	+0.24	+0.24	+0.24
Gypsum	+0.07	+0.07	+0.07	+0.07
Subtotal	**−4.39**	**−4.65**	**−4.68**	**−2.90**

The aggregate non-priority material contribution remained materially smaller than the aluminium-route results emphasized in the main text. Additionally, the non-priority results were strongly shaped by a small number of material families. In all scenarios, the wood-based components represented by the CLT and SPF softwood proxies contributed the largest negative net terms, reflecting the use of energy-recovery proxy treatments with substantial embedded avoided-burden effects. In the contemporary case, these terms were −5.81 kg CO_2e/m^2 for CLT and −1.50 kg CO_2e/m^2 for SPF softwood; comparable values were observed in the prospective cases, with somewhat smaller absolute credits in SSP5-H. By contrast, EPDM was consistently the largest positive contributor, with a net burden of approximately +2.25 kg CO_2e/m^2 in the contemporary case and +2.06 to +2.21 kg CO_2e/m^2 across the prospective scenarios. The remaining non-priority materials, including polyamide, glass, silicone, silica gel, XPS, polyethylene, and gypsum, were individually small, generally remaining well below ±0.25 kg CO_2e/m^2 each, with the exception of gypsum's modest positive contribution of roughly +0.07 kg CO_2e/m^2. Argon was assumed to be released directly to the biosphere at end of life and therefore did not contribute a distinct GWP100 burden in the foreground treatment results. The metallic coating was treated as negligible in mass and was assumed to be embedded within the insulated glazing unit representation rather than modelled as a separate EoL flow.

These results support the conclusion that the non-priority materials were treated consistently and do not reveal a hidden omitted burden large enough to overturn the aluminium-focused interpretation. This pattern is also methodologically plausible in a consequential framework: where end-of-life routes broadly reflect the prevailing treatment logic already embedded in the background system, rerouting the remaining façade materials often produces only limited net change unless the route activates a strong substitution claim. In the present case, the main departure from that pattern is the wood layer, where the large negative net values are driven by the selected energy-recovery proxies rather than by a second independent hotspot comparable to aluminium. In ecoinvent, waste-treatment activities are not modelled as pure disposal sinks. They include direct emissions, material and energy inputs, and, where relevant, by-products such as electricity, heat, or secondary materials. ecoinvent also explicitly models landfill gas capture and utilisation in some landfill datasets, which means that small non-zero Stage D terms can legitimately appear even for disposal-oriented proxy routes. Once those embedded exchanges are split into explicit C3–C4 and Stage D wrapper nodes and relinked to localized providers, the resulting Stage D term may remain small and may be either slightly positive or slightly negative depending on the product exchanged and the impact category being assessed. More generally, landfill gas recovery is known to reduce burdens in some indicators while increasing them in others, so a small positive residual associated with landfill-based treatment is not inherently anomalous.

The large credit associated with the mass-timber proxies should therefore be interpreted primarily as a modelling consequence of the ecoinvent treatment representation used here, namely that energy recovery is treated as a consequential departure from disposal rather than as a neutral background norm. That treatment choice was retained because it better reflects the intended Canadian end-of-life direction for wood, but the non-priority sectors were not manually overridden to a fully Canada-specific treatment mix, and this remains a data limitation. Overall,

the non-priority total is partially credit-dominated, but in a much narrower and less decision-relevant way than the aluminium routes. In practice, the aggregate negative total is driven mainly by the selected wood energy-recovery proxies, while most of the remaining material families behave as small positive residual burdens.

Appendix C — Background Database Construction, QA, and Reproducibility

C.1 Purpose & Scope

This appendix documents the creation, freezing, and quality assurance of the contemporary and prospective background databases used in the thesis. The purpose of this appendix is to provide the minimum reproducibility and audit information needed to support the results without reproducing every raw log file, checksum file, or database manifest in the thesis itself.

Appendix C reports:

1. The version and configuration statement for the background systems,
2. the freeze record for the contemporary baseline background,
3. the freeze records for the prospective scenario backgrounds,
4. the summary QA outcomes that establish their internal consistency,
5. and the reproducibility note that links the thesis to the archived build artifacts in the repository.

This appendix should be read together with Appendix B, which documents the foreground route models layered on top of these backgrounds, and Appendix D, which explains which market constructs were inherited directly from the scenario databases and which remained foreground defined.

C.2 System and Version Statement

The thesis uses ecoinvent 3.10.1 consequential as the base background system for the contemporary case. Prospective backgrounds were generated using premise 2.3.2 coupled to IMAGE-based scenario information for the 2050 assessment year under the selected scenario set:

- ***SSP1-VLLO***
- ***SSP2-M***
- ***SSP5-H***

All prospective backgrounds were treated as single-year 2050 snapshots under the perfect foresight (PERF) configuration. No rolling-year averaging, multi-year blending, or post-build manual restructuring was retained in the final background databases used for the thesis.

The environment recorded in the freeze records below was:

- Python 3.11.14
- Windows-10-10.0.26200-SP0
- bw2data 4.5.3

- bw2io 0.9.11
- bw2calc 2.2.2
- premise 2.3.2

C.3 Contemporary Background Snapshot

C.3.1 Role of the Contemporary Snapshot

The contemporary background provides the fixed reference world against which all prospective comparisons are made. It was not generated through premise transformation. Rather, it corresponds to the frozen consequential database used as the baseline background for the case study.

C.3.2 Contemporary Freeze Record

Table C-1 reports the metadata for the contemporary project.

Table C-1*: Contemporary background snapshot record*

Item	***Value***
Snapshot ID	S-BASE-2025-10-10
Timestamp (UTC)	2026-03-10T21:09:06.177202+00:00
Brightway project	pCLCA_CA_2025_contemp
Database name	ecoinvent_3.10.1.1_consequential_unitprocess
Base role	Frozen contemporary baseline background
Python	3.11.14
Platform	Windows-10-10.0.26200-SP0
bw2data	4.5.3
bw2io	0.9.11
bw2calc	2.2.2
premise	2.3.2
Activities	20,772
Total exchanges	610,085
Technosphere exchanges	254,178
Biosphere exchanges	355,907

Production exchanges	20,772
Substitution exchanges	0
Unique locations	330
Activities missing location	0
Missing input references	0
Missing input database references	0
Broken internal technosphere links	0
Unexpected external technosphere dependencies	0
Biosphere databases detected	ecoinvent-3.10-biosphere
Matrix processable	TRUE
Matrix process message	ok
Label hash	d065e11abbeeb4119849ef7dff86112a8b5a3d32
Record SHA-256	16b38746bebe1c2a5c7c677789057a1b47f5f279c3117e89b4504298f89cfb54

C.3.3 Contemporary Archived Artefacts

The contemporary snapshot is supported by archive-level checksums that make the frozen baseline auditable without reproducing the underlying files in the thesis itself Table C-2.

Table C-2: *Contemporary archived artifacts and SHA-256 hashes*

Artifact	***Recorded path***	***SHA-256***
Snapshot zip	C:\brightway_workspace\snapshots\S-BASE-2025-10-10_ecoinvent-3.10.1-consequential.zip	7e3550b6bc3a141517acb79d3274acbe67f594e967a52e4bedc4e43792adba8c
Snapshot manifest	C:\brightway_workspace\snapshots\S-BASE-2025-10-10_ecoinvent-3.10.1-consequential\S-BASE-2025-10-10_manifest.json	9a2bb715fc5edc99b1c4de13472ef279482a44f7e57b2e37d1acd2fa31c73967

Processed zip	C:\brightway_workspace\snapshots\S-BASE-2025-10-10_ecoinvent-3.10.1-consequential\ecoinvent-3.10.1-consequential.processed.zip	4ba325d08d259e48f2e16b6461af6e70e915c2c8f2e4fd0f611006a37f53c495

C.3.4 Interpretation Note

The contemporary background was treated as frozen for all downstream modelling. Later foreground route wrappers, marginal-supplier constructs, and uncertainty runs therefore referenced a stable baseline background rather than a changing or partially rebuilt base system.

C.4 Prospective Scenario Backgrounds Records

C.4.1 Role of the Prospective Set

The prospective backgrounds define three internally consistent 2050 scenario worlds used to evaluate the aluminium end-of-life pathways under alternative future energy-system conditions. Their role is not to represent a fully rebuilt 2050 economy across all sectors, but rather to provide coherent pathway-specific future backgrounds onto which the foreground route models can be layered.

C.4.2 Prospective Fingerprint Records

Table C-3 reports the structural and processability checks for the frozen prospective backgrounds, while Table C-4 records the scenario-specific identifiers used to distinguish and audit those databases.

***Table C-3**: Prospective scenario background structural QA summary*

Metric	***SSP1-VLLO***	***SSP2-M***	***SSP5-H***
Database name	prospective_conseq_IMAGE_SSP1VLLO_2050_PERF	prospective_conseq_IMAGE_SSP2M_2050_PERF	prospective_conseq_IMAGE_SSP5H_2050_PERF
Timestamp (UTC)	2026-03-10T21:15:00.682128+00:00	2026-03-10T21:15:00.682128+00:00	2026-03-10T21:15:00.682128+00:00
Brightway project	pCLCA_CA_2025_prospective	pCLCA_CA_2025_prospective	pCLCA_CA_2025_prospective
Activities	26,665	26,664	26,664
Total exchanges	734,784	734,393	734,519
Technosphere exchanges	313,406	313,020	313,156
Biosphere exchanges	421,378	421,373	421,363
Production exchanges	26,665	26,664	26,664
Substitution exchanges	0	0	0
Unique locations	356	356	356
Activities missing location	0	0	0
Missing input references	0	0	0
Missing input database references	0	0	0
Broken internal technosphere links	0	0	0
Unexpected external technosphere dependencies	0	0	0
Biosphere databases detected	ecoinvent-3.10.1-biosphere	ecoinvent-3.10.1-biosphere	ecoinvent-3.10.1-biosphere
Matrix processable	TRUE	TRUE	TRUE
Matrix process message	ok	ok	ok

***Table C-4**: Prospective scenario background identifiers and integrity records*

Identifier / checksum field	**SSP1-VLLO**	**SSP2-M**	**SSP5-H**
Database name	prospective_conseq_IMAGE_SSP1VLLO_2050_PERF	prospective_conseq_IMAGE_SSP2M_2050_PERF	prospective_conseq_IMAGE_SSP5H_2050_PERF
Label hash	d1e95bbb4269eb7abe65a04c6c3883b7da2aa9e7	e5cdd41fb06b7911f904c752c33e7560d612b558	17ba3602c83fcee67f49a0531893a923b00ac39a
Record SHA-256	280823b06be67a5117749a01d5ddca557461b1fd53187ea49deca434cd86fdc0	bafbc93e723c1f6f2c0c713b851f81f6cbb79351acfa69916085d8cb45508111	d3d9c3e751749d2e0f713f035a9d019145752fcc0f43bac99bd6c108f86df688

C.4.3 Interpretation Note

These records confirm that the prospective databases used in the thesis were present, processable, and internally coherent at the point of audit. The thesis treats these databases as the default frozen scenario backgrounds used in all downstream prospective modelling.

C.4.4 Scope of Prospective Transformation

Only the electricity system and the fuel and energy-carrier supply system were structurally transformed in the prospective backgrounds. Other sectors, including cement, steel, freight transport, industrial heat, and direct air capture, were not fully rewritten into 2050 structures. These sectors therefore operate under the transformed future energy background while retaining the structural form inherited from the base consequential system.

This energy-first transformation should be interpreted as a deliberate scope choice rather than an omission. It prioritizes internal scenario consistency for the most consequential background drivers while avoiding partially rebuilt future sectors that could not be updated consistently at the required consequential resolution.

C.5 Build Procedure Summary

C.5.1 Contemporary Background

The contemporary background required no premise transformation. Its preparation therefore consisted primarily of freezing and documenting the base consequential system used in the case study.

C.5.2 Prospective Backgrounds

Each prospective background followed the same general deterministic sequence:

(1) selection of the IMAGE pathway and 2050 assessment year,
(2) generation of pathway-year electricity and fuel / energy-carrier markets through premise,
(3) writing of the transformed database to the Brightway environment,
(4) execution of deterministic inspection and audit checks,
(5) recording of fingerprint and manifest artifacts,
(6) designation of the resulting database as the frozen scenario-year background for downstream use.

No manual bridge-market creation or ad hoc reweighting was retained as part of the final frozen scenario backgrounds.

C.6 Quality Assurance Summary

C.6.1 Purpose of Quality Assurance

The quality assurance (QA) stage was designed to confirm that the background databases were internally consistent and computationally valid, not to reopen or reinterpret the pathway translation itself.

C.6.2 QA Checks Applied

Table C-5 details the QA checks applied to the databases used in this analysis as well as the rationale.

Table C-5: *Quality-assurance checks applied*

QA family	***Check***	***Purpose***
Location completeness	Activities missing location = 0	Confirms geography fields are populated for later routing and inspection
Link completeness	Missing input references = 0	Confirms exchanges do not point to unresolved inputs
Database completeness	Missing input database references = 0	Confirms exchanges do not point to absent databases

Internal technosphere integrity	Broken internal technosphere links = 0	Confirms no broken internal foreground/background links
External dependency scan	Unexpected external technosphere dependencies = 0	Confirms no unintended dependencies outside the intended database set
Biosphere consistency	Single expected biosphere database detected	Confirms biosphere links are resolved to the intended biosphere system
Matrix processability	Matrix processable = TRUE	Confirms Brightway can process the database without error

Across the contemporary and prospective databases, all audited QA fields returned a pass-equivalent outcome: no missing input references, no missing input database references, no broken internal technosphere links, no unexpected external technosphere dependencies, no missing activity locations, and successful matrix processing in all cases. On this basis, all four background databases were treated as computationally valid and suitable for downstream use.

C.7 Limitations and Interpretation

Appendix C establishes that the backgrounds used in the thesis were frozen, auditable, and computationally valid. It does not imply that the prospective backgrounds represent a fully rebuilt future economy across all sectors. Rather, they should be interpreted as scenario-consistent 2050 backgrounds in which the principal structural transformation occurs through electricity and fuel/energy-carrier markets, with other sectors inheriting those conditions while retaining their base structural configuration.

Appendix D — Markets, Marginal Suppliers, and Data Provenance

D.1 Purpose & Scope

This appendix documents the market-definition and marginal-supplier logic used to support consequential linking and Stage D substitution in the case study. It records the displaced product markets adopted for the retained end-of-life routes, the evidence hierarchy used to identify marginal suppliers, the principal assumptions used to delimit geography and time horizon, and the summarized marginal mix outputs used in the deterministic and uncertainty analyses.

The purpose of this appendix is to provide a transparent audit trail. It is not intended to reproduce every intermediate data table, trade-filter worksheet, growth calculation, or scenario extraction in full. Those materials, including raw source sheets, HS-code filters, detailed $V_0/V_1/g/r/\tau/\Delta C/\alpha$calculations, structural-break notes, and scenario extract files, are archived in the accompanying code repository.

This appendix should be read together with Appendix A, which documents Stage D claim governance; Appendix B, which documents the foreground route models; Appendix C, which documents the frozen background databases; and Appendix F, which documents uncertainty and sensitivity cases that bear on market interpretation.

D.2 Market Boundary Register

D.2.1 Role of Market Definition in Case Study

Stage D substitution depends not only on the quantity and quality of a recovered output, but also on the definition of the market into which that output is assumed to enter. In this thesis, market boundaries were defined before crediting and were based on:

1. functional equivalence,
2. grade or specification compatibility where relevant,
3. consequential geography and trade exposure.

The purpose of the market boundary register (Table D-1) is to state, for each retained route, the largest defensible displaced market that the recovered output could plausibly serve.

***Table D-1**: Market boundary register for retained displaced products and services*

Route / credited output	***Displaced market***	***Functional / specification basis***	***Geography***	***Marginal mix treatment***
Reuse	Aluminium non-structural façade components	Like-for-like façade-component service; fit, geometry, and service duty govern eligibility	Canada-wide	Final implementation retained the façade-component/service claim, but Stage D was represented using the same primary-aluminium marginal mix as conventional recycling together with an avoided-extrusion proxy
Conventional recycling	Wrought extrusion feedstock (billet)	Billet suitable for 6000-series extrusion after standard cast-house conditioning	Canada-wide	Foreground-constructed for contemporary; prospective supplier-set overlay
MS-FSC	Wrought extrusion feedstock (billet)	Billet-equivalent output suitable for 6000-series extrusion, subject to pass / quality gate	Canada-wide	Same displaced market class as conventional recycling
Hydrolysis — H_2	Industrial merchant gaseous hydrogen	Bulk industrial gaseous hydrogen, independent of production route	Canada-wide	Contemporary foreground-constructed hydrogen marginal mix; final central case resolved to Alberta Energy Regulator SMR. Prospective hydrogen inherited where represented
Hydrolysis — $Al(OH)_3$	Aluminium hydroxide supply market	Routine solids-handling compatible aluminium hydroxide product form	Background market geography	Inherited from background market; no custom marginal mix constructed

			used in implementation	
Landfill	None	No displaced product or service	N/A	No Stage D credit
Route / credited output	Displaced market	Functional / specification basis	Geography	Marginal mix treatment
Reuse	Aluminium non-structural façade components	Like-for-like façade-component service; fit, geometry, and service duty govern eligibility	Canada-wide	Final implementation retained the façade-component/service claim, but Stage D was represented using the same primary-aluminium marginal mix as conventional recycling together with an avoided-extrusion proxy

D.3 Source Hierarchy and Data Provenance

D.3.1 General Hierarchy

Marginal supplier identification did not rely on a single universal data source. Instead, a hierarchy was applied depending on market type, time horizon, and whether the market was treated as a foreground construct or inherited directly from the scenario database.

The hierarchy used in the thesis was:

1. Scenario-database inherited commodity mix,
2. Canada-anchored forward-looking projection or outlook,
3. Observed Canadian production plus reconciled trade evidence,
4. Heuristic supplier-set–constrained construction.

Table D-2 details which route-scenario combination relied on which evidence source.

***Table D-2**: Evidence families used to define markets and marginal suppliers*

Market	***Primary evidence family***	***Secondary support***	***Final treatment***
Contemporary primary aluminium billet	Canadian facility-resolved production + trade-screened import evidence	Industry and technology context	Foreground-constructed
Contemporary extrusion / reuse proxy	Canadian extrusion production + screened imports	Sector reviews and reuse practice literature	Foreground-constructed proxy
Contemporary hydrogen	Public project and production disclosures aggregated to technology–region bins	Hydrogen strategies, regulatory and industry summaries	Foreground-constructed heuristic
Contemporary electricity	Forward-looking Canadian generation projections (2025–2030)	Technology-level planning evidence	Foreground-constructed
Prospective electricity and fuels	Scenario database commodity mixes	Scenario documentation	Inherited directly
Prospective hydrogen	Scenario database where represented	Scenario documentation	Inherited directly
Prospective primary aluminium	Contemporary supplier-set retained; technology overlay applied	Long-horizon technology-change evidence	Hybrid foreground overlay

D.3.2 Repository Note

The detailed source crosswalks, including HS-code filters, bilateral partner mappings, unit conversions, projection harmonization notes, and raw source tables, are stored in the repository and are not reproduced in full here.

D.4 Marginal Identification Rules & Trend Detection Logic

The marginal mix development followed the rules outline in Chapter 3. However, the full structure denoted there was not applied in every market due to data limitations and/or unique cases. The structurally variations are as follows:

- Contemporary electricity used a forecast-based net-change normalization over 2025–2030 and did not apply capital-replacement correction.
- Contemporary primary aluminium used a fixed 2018–2023 historical window as a documented exception.
- Contemporary hydrogen used a multi-year realized-output screen where possible but was ultimately treated as a bounded heuristic.
- Prospective electricity, fuels, and hydrogen were inherited directly from the scenario databases and were not recomputed using the historical equations.
- Prospective primary aluminium retained the contemporary supplier set while overlaying technological evolution.

D.5 Contemporary Foreground-Constructed Marginal Mixes

D.5.1 Primary Aluminium Billet Market

The contemporary primary aluminium billet marginal mix was derived from a facility-resolved domestic production screen over 2018–2023, followed by an import-enabled check using trade-screened billet-equivalent unwrought aluminium. Imports were retained as an explicit test but did not materially affect the central case. The resulting contemporary marginal mix is therefore dominated by expanding domestic facilities (Table D-3).

Table D-3: *Contemporary marginal supplier mix for the aluminium billet market (Canada; 2018–2023 basis)*

Supplier	***Annualized growth g_i(%/yr)***	***Positive annual ΔV(kt/yr)***	***Included in marginal set?***	***Marginal mix (%)***
Rio Tinto – Alma Smelter	0.82	3.80	Yes	15.45
Aluminerie Alouette – Alouette Smelter	1.71	10.00	Yes	40.65

Alcoa – Deschambault Smelter	2.08	5.40	Yes	21.95
Rio Tinto – Grande Baie Smelter	-0.34	0.00	No	0.00
Alcoa – Bécancour Smelter	1.23	5.40	Yes	21.95
Rio Tinto – Kitimat Smelter	-2.71	0.00	No	0.00
Total		24.60		100.00

This mix should be interpreted as a foreground consequential approximation of which suppliers were adding effective capacity, not as a production-share table. The detailed facility worksheets, structural-break notes, and import-enabled sensitivity cases are archived in the repository.

D.5.2 Reuse Proxy/Extrusion Market

The reuse pathway was linked to a Canada-wide displaced market for aluminium façade-component service. However, a directly observable national time series for reusable façade-component supply was not available. For that reason, the practical evidence used to bound marginal supply was informed by domestic aluminium semi-fabrication and screened import evidence rather than by a dedicated reuse-market dataset.

This distinction is important. The market boundary for reuse remained a façade-component service market. The evidence proxy used to bound marginal supply was drawn from the closest measurable aluminium supply chain evidence available. This avoids collapsing reuse into the same displaced market as billet recycling while still grounding the central case in observable supply data. In the thesis, it is sufficient to record that:

- the reuse route displaces a façade-component service market,
- the evidence used to bound marginal supply was informed by domestic semi-fabrication / extrusion evidence and screened imports,
- and the full supplier worksheets are retained outside the appendix.

D.5.3 Hydrogen Market

The contemporary industrial hydrogen market was constructed as a bounded heuristic because project and production data were sparse, heterogeneous, and unevenly resolved across technologies and regions. After aggregation into technology–region families and application of the expansion-signal screen, the only supplier configuration retained in the central case was Alberta conventional fossil-based hydrogen (SMR / ATR without CCS).

This mix should be read as a conservative central case rather than a claim that low-carbon hydrogen is irrelevant.

D.5.4 Electricity Market

Electricity was not treated as a Stage D product in this thesis, but it was a major burden-side driver for the foreground route models. A Canada-wide marginal electricity mix was therefore constructed for contemporary foreground operations using forward-looking CER generation projections over 2025–2030 (Table D-4). The CER's Energy Futures platform is specifically intended to explore long-term possible energy futures and provides downloadable projection data and interactive visualization tools for that purpose.

Table D-4*: Canada-wide marginal electricity mix derived from 2025–2030 projected generation changes*

Supplier	***GWh (2025)***	***GWh (2030)***	***ΔV(GWh)***	***Positive ΔV^{+}(GWh)***	***Marginal share (%)***
Hydro / Wave / Tidal	403,350.84	437,850.35	34,499.51	34,499.51	40.09
Wind	57,101.81	80,824.64	23,722.83	23,722.83	27.56
Natural Gas	106,529.07	128,015.91	21,486.84	21,486.84	24.97
Uranium	78,631.37	82,253.63	3,622.26	3,622.26	4.21
Biomass / Geothermal	7,838.55	9,204.80	1,366.25	1,366.25	1.59
Solar	12,140.31	13,504.24	1,363.93	1,363.93	1.58
Coal & Coke	8,132.43	654.81	-7,477.62	0.00	0.00
Oil	1,379.91	683.79	-696.12	0.00	0.00
Total	675,104.29	752,992.17	77,887.88	86,061.62	100.00

Province-specific marginal electricity variants were also computed where foreground alignment required them, particularly for route contexts tied to province-resolved burden-side electricity assumptions. Province-specific electricity variants are provided in the repository.

D.6 Prospective Inherited Commodity Mixes and Foreground Overlays

D.6.1 Inherited Prospective Commodity Mixes

For the 2050 prospective case, the thesis does not recompute electricity, fuels, and hydrogen marginality using the historical foreground equations. Where the scenario database already provides the relevant pathway-year commodity representation, those mixes are treated as authoritative scenario inputs.

This applies to:

- electricity,
- fuel and energy carriers,
- hydrogen where represented.

D.6.2 Primary Aluminium Overlay

Primary aluminium was treated differently from electricity, fuels, and hydrogen. The thesis did not impose plant-level 2050 forecasts for Canadian aluminium suppliers. Instead, it retained the contemporary supplier set as a "who supplies?" construct and applied a technology-evolution overlay to represent changes in how that supply is produced under 2050 conditions. This choice was adopted for two reasons.

First, a direct plant-by-plant 2050 supplier forecast was not available at a level of resolution that could be defended consistently within the study's consequential market logic. In contrast, the premise workflow explicitly integrates IAM variables for electricity and for fuels and hydrogen supply chains, so those sectors can be inherited more directly from the scenario database. The premise documentation states that electricity generation and fuels and hydrogen supply chains are among the subset of IAM variables currently integrated, whereas full consistency across all sectors is not yet achieved (Sacchi et al., 2022).

Second, there is credible Canadian evidence that the production mechanism of marginal aluminium supply may change over the study horizon even if the identity of future suppliers cannot be forecast robustly. Industry and technology sources document both continuing expansion of low-carbon aluminium capacity in Canada and active commercialization of inert-anode technology. The Aluminium Association of Canada records major recent investments in Canadian low-carbon aluminium and billet capacity, including Rio Tinto's AP60 expansion and Alma billet-capacity additions, while ELYSIS documents the progression from Canadian industrial demonstration to commercialization of inert-anode technology (NRCan, 2018; Rio Tinto, 2023)

On that basis, the overlay was used to modify the "how is supply produced?" dimension without inventing a plant-resolved 2050 market-share forecast. At minimum, the overlay distinguishes:

- conventional supply,
- inert-anode / ELYSIS-type supply.

A similar forward-looking expansion narrative exists for hydrogen, but hydrogen is already represented within the IAM–premise fuels framework and was therefore inherited directly where represented rather than being rebuilt manually as a separate foreground overlay. The detailed implementation, scenario-specific provider mapping, and sensitivity variants for the aluminium overlay remain documented in the supporting materials.

D.7 Limitations

Appendix D is intentionally summary-oriented, but several limitations should be stated explicitly.

First, the reuse route uses a market boundary / evidence-proxy split. The displaced market is façade-component service, but the practical supplier evidence used to bound that market is drawn from the closest measurable aluminium semi-fabrication evidence. This is defensible, but it is not the same as having a directly observable national reuse-supply time series.

Second, the contemporary hydrogen market is a bounded heuristic. The central case is deliberately conservative and reflects the strongest defensible supplier family under the available evidence. It should not be overread as a claim that Alberta conventional hydrogen is the only plausible future responder.

Third, the prospective primary aluminium treatment is a hybrid foreground overlay. It is more explicit and auditable than silently inheriting a market that is poorly represented for the Canadian decision context, but it is still a modelling construct rather than a plant-by-plant forecast.

Fourth, the contemporary electricity market is constructed from projection-based generation change, not from observed historical marginality. This is a deliberate choice consistent with the electricity special case and with the planning-oriented role electricity plays in the foreground route burdens, but it remains a modelling choice and should be interpreted as such. The CER sources establish the relevance of Canada-wide interconnection and provide the projection platform from which the electricity values were extracted.

Appendix E — Additional LCIA Categories and Qualitative Stability

E.1 Purpose & Scope

This appendix reports the cross-indicator screening used to test whether the qualitative interpretation inferred from the primary climate-change indicator remained stable across the broader ReCiPe 2016 Midpoint (H) method set. The purpose of this appendix is not to reproduce every route-by-scenario result in full, but rather to document additional evidence needed to support or qualify the GWP-based interpretation in the main text.

Three questions guide this appendix section:

(7) To what extent does the net sign observed under GWP100 remain aligned across the non-climate midpoint indicators?
(8) To what extent does the ranking of routes under GWP100 persist when other midpoint mechanisms are considered?
(9) Where qualitative divergence occurs, does it appear to arise primarily from the burden-side treatment chain (C3–C4) or from the receiving-system credit term (Stage D)?

The detailed method-by-method result tables, plotting scripts, full heatmap inputs, and all supplementary decomposition figures are archived in the accompanying code repository. This appendix reports only the summary diagnostics required for interpretation.

E.2 Screening Dataset & Indicator Set

The screening was performed using the ReCiPe 2016 Midpoint (H) method set computed for the five assessed end-of-life routes across the four scenario contexts:

- ***Contemporary,***
- ***SSP1-VLLO,***
- ***SSP2-M,***
- ***SSP5-H.***

The evaluated route set comprised:

- ***landfill,***
- ***reuse,***
- ***conventional recycling,***
- ***aluminium hydrolysis,***
- ***MS-FSC.***

The uploaded midpoint-method index contains 18 ReCiPe 2016 Midpoint (H) indicators. Of these, GWP100 was treated as the primary reference indicator used in the main text and the remaining 17 indicators were treated as the non-climate midpoint set for sign-consistency screening. Table E-1 summarizes the resultant indicators.

***Table E-1**: ReCiPe 2016 Midpoint (H) indicators included in the screening set*

Group	***Indicator***
Primary climate reference	climate change | global warming potential (GWP100)
Non-climate midpoints	acidification: terrestrial | terrestrial acidification potential (TAP)
	ecotoxicity: freshwater | freshwater ecotoxicity potential (FETP)
	ecotoxicity: marine | marine ecotoxicity potential (METP)
	ecotoxicity: terrestrial | terrestrial ecotoxicity potential (TETP)
	energy resources: non-renewable, fossil | fossil fuel potential (FFP)
	eutrophication: freshwater | freshwater eutrophication potential (FEP)
	eutrophication: marine | marine eutrophication potential (MEP)
	human toxicity: carcinogenic | human toxicity potential (HTPc)
	human toxicity: non-carcinogenic | human toxicity potential (HTPnc)
	ionising radiation | ionising radiation potential (IRP)
	land use | agricultural land occupation (LOP)
	material resources: metals/minerals | surplus ore potential (SOP)
	ozone depletion | ozone depletion potential (ODP)
	particulate matter formation | particulate matter formation potential (PMFP)
	photochemical oxidant formation: human health | photochemical oxidant formation potential: humans (HOFP)
	photochemical oxidant formation: terrestrial ecosystems | photochemical oxidant formation potential: ecosystems (EOFP)
	water use | water consumption potential (WCP)

All results were evaluated using the same decomposition structure used in the main text:

- ***C3–C4*** burdens,
- ***Stage D*** credit term,
- ***net*** = C3–C4 + Stage D.

The sign-consistency screening uses the net result as the primary comparison basis, while representative decomposition figures are used to localize where divergence arises.

E.3 Reporting Convention & interpretation Rule

The purpose of this appendix is not to ask whether every indicator agrees numerically with GWP100, but whether the qualitative interpretation suggested by GWP100 remains defensible when other midpoint mechanisms are considered.

Two diagnostics are therefore used:

E.3.1 Sign Consistency

For each route–scenario cell, the sign-consistency fraction reports the share of the 17 non-climate midpoint indicators whose net sign aligned with the corresponding GWP100 net sign for that same cell. Thus:

- **17/17** indicates full qualitative sign alignment,
- lower values indicate that some non-climate midpoint indicators imply an opposite net direction relative to GWP100.

E.3.2 Rank Inversion

For each midpoint indicator and scenario, the five routes were ordered by their net result. A rank inversion was recorded when that route ordering differed from the ordering obtained under GWP100 for the same scenario.

This is a stricter diagnostic than sign consistency. An indicator can agree with GWP100 on sign while still changing the relative ordering among reuse, recycling, hydrolysis, and MS-FSC. Rank inversion therefore captures instability in comparative interpretation even when the direction of benefit or burden does not reverse.

E.4 Sign-Consistency Screening Against GWP100

E.4.1 Route-by-Scenario Sign-Consistency Matrix

Table E-2 reports the sign-consistency fractions for each route–scenario cell aligned with the sign-consistency heatmap summary in Figure 4-11.

***Table E-2**: Sign consistency versus GWP100 across the 17 non-climate midpoint indicators*

Route	***Contemporary***	***SSP1-VLLO***	***SSP2-M***	***SSP5-H***
Landfill	17/17 (1.000)	16/17 (0.941)	17/17 (1.000)	16/17 (0.941)
Reuse	16/17 (0.941)	16/17 (0.941)	17/17 (1.000)	16/17 (0.941)
Conventional recycling	15/17 (0.882)	16/17 (0.941)	16/17 (0.941)	16/17 (0.941)
Aluminium hydrolysis	14/17 (0.824)	12/17 (0.706)	12/17 (0.706)	12/17 (0.706)
MS-FSC	14/17 (0.824)	15/17 (0.882)	15/17 (0.882)	14/17 (0.824)

This matrix supports four immediate observations. First, landfill is the most sign-stable route overall. Its sign aligns with GWP100 in either all or nearly all non-climate indicators across scenarios. This is expected because landfill has no Stage D credit and therefore behaves as a burden-only baseline. Second, reuse is also highly sign stable. It retains 16/17 or 17/17 alignment in all four scenarios, indicating that GWP100 is a reasonably reliable directional proxy for this route, even though it does not guarantee stable route ranking. Third, conventional recycling remains comparatively stable on sign but less so than reuse. Its lowest alignment occurs in the contemporary case, where 15 of 17 non-climate indicators align with GWP100. Fourth, the lowest sign-consistency is concentrated in the emerging routes, especially aluminium hydrolysis, which falls to 12/17 aligned indicators in all three 2050 scenarios and only reaches 14/17 in the contemporary case. MS-FSC is more stable than hydrolysis, but still exhibits lower sign-consistency than landfill, reuse, or conventional recycling.

E.4.2 Route-Level Average Sign Consistency

Averaged across the four scenarios, the sign-consistency fractions are:

- Landfill: **0.971**
- Reuse: **0.956**
- Conventional recycling: **0.926**
- MS-FSC: **0.853**
- Aluminium hydrolysis: **0.735**

These averages confirm that the routes most dependent on more complex Stage D structures, especially cross-sector or higher-uncertainty receiving systems, are also those for which GWP100 is the least reliable qualitative proxy across the midpoint set.

E.5 Indicators Exhibiting Sign Reversals

A sign reversal was counted when at least one route–scenario cell had a net sign opposite to that observed under GWP100 for the same cell. 9 of the 17 non-climate midpoint indicators exhibited at least one sign reversal (Table E-3).

Table E-3: *Non-climate midpoint indicators with at least one sign reversal relative to GWP100*

Indicator	***Cells with sign reversal (out of 20 route–scenario cells)***	***Primary concentration of reversals***
Ionising radiation (IRP)	9	Mixed; includes multiple contemporary metal-displacement routes and selected landfill/reuse/MS-FSC cells in SSP1-VLLO and SSP5-H
Freshwater ecotoxicity (FETP)	7	Prospective hydrolysis and MS-FSC; contemporary MS-FSC
Marine ecotoxicity (METP)	7	Prospective hydrolysis and MS-FSC; contemporary MS-FSC
Terrestrial ecotoxicity (TETP)	4	Conventional recycling in all four scenarios
Freshwater eutrophication (FEP)	3	Hydrolysis in all three 2050 scenarios
Marine eutrophication (MEP)	3	Hydrolysis in all three 2050 scenarios
Ozone depletion (ODP)	3	Hydrolysis in all three 2050 scenarios
Human toxicity, carcinogenic (HTPc)	1	Contemporary hydrolysis only
Surplus ore potential (SOP)	1	Contemporary hydrolysis only

This distribution is important for interpretation.

The largest sign instability appears in ionising radiation, followed by freshwater and marine ecotoxicity, with additional route-specific reversals appearing in terrestrial ecotoxicity, freshwater eutrophication, marine eutrophication, ozone depletion, and isolated hydrolysis reversals in HTPc and SOP.

By contrast, the following non-climate indicators showed no sign reversals relative to GWP100 in any of the 20 cells:

- terrestrial acidification potential (TAP),
- fossil fuel potential (FFP),
- human toxicity potential (HTPnc),
- agricultural land occupation (LOP),
- particulate matter formation potential (PMFP),
- photochemical oxidant formation potential: humans (HOFP),
- photochemical oxidant formation potential: ecosystems (EOFP),
- water consumption potential (WCP).

This does not mean that these indicators are fully stable in comparative terms; it means only that they do not reverse the net sign observed under GWP100.

E.6 Rank-Inversion Diagnostics

Sign stability alone is insufficient to guarantee that the route ordering observed under GWP100 is preserved. The rank-inversion diagnostic therefore compares the route ordering under each midpoint indicator with the route ordering under GWP100 within the same scenario.

E.6.1 Summary Result

16 of the 17 non-climate midpoint indicators exhibited at least one rank inversion relative to GWP100. Only FFP preserved the same route ordering as GWP100 across all four scenarios. This is detailed in Table E-4.

Table E-4*: Rank-inversion summary relative to GWP100*

Number of cenarios with at least one rank inversion	***Number of non-climate indicators***	***Indicators***
4 of 4 scenarios	11	TAP, FETP, METP, TETP, FEP, HTPnc, IRP, LOP, ODP, PMFP, HOFP
3 of 4 scenarios	3	MEP, EOFP, WCP
1 of 4 scenarios	2	HTPc, SOP

0 of 4 scenarios	1	FFP

The rank-inversion results support several conclusions. First, comparative route ranking is far less stable than net sign. This is especially important where the main text draws distinctions among reuse, conventional recycling, hydrolysis, and MS-FSC. Second, emerging routes are not the only drivers of instability. Some indicators, such as TAP and HTPnc, preserve sign alignment but still reorder routes in every scenario. Third, FFP is the only non-climate indicator that remains fully ranking-stable relative to GWP100 across all scenarios. This suggests that fossil-energy demand tracks the same broad comparative signal as climate change more closely than most other midpoints in this case study.

E.7 Representative Indicators and Qualitative Localization of Instability

This section provides brief narrative interpretation for the representative decomposition figures used to localize where instability arises.

E.7.1 Human Toxicity: Carcinogenic (HTPc)

HTPc is an example of a localized and route-specific sign reversal, rather than a widespread instability. Only one route–scenario cell reversed sign relative to GWP100: contemporary hydrolysis.

This makes HTPc useful as a bounded example. It shows that even when an indicator is not globally unstable, isolated route-specific divergences can still matter for interpretation. In this case, the divergence is concentrated in a route that already exhibits strong receiving-system dependence and relatively high burden-side sensitivity.

E.7.2 Acidification: Terrestrial (TAP)

TAP is an example of an indicator with no sign reversals but full ranking instability. Across all four scenarios, the route ordering differs from GWP100 even though the net sign generally remains directionally aligned.

This distinction is important because it shows that GWP100 can remain a valid directional signal while still failing to preserve comparative route ranking. In other words, TAP does not contradict the broad climate interpretation on sign, but it does change which route looks preferable within the crediting set.

E.7.3 Freshwater and Marine Ecotoxicity (FETP and METP)

FETP and METP are the clearest examples of burden-shifting instability in the uploaded decomposition figures. Their sign reversals are concentrated in:

- prospective hydrolysis,
- prospective MS-FSC,
- and contemporary MS-FSC.

These ecotoxicity indicators therefore behave differently from GWP100 because the burden and receiving-system structures are environmentally different under ecotoxicity mechanisms. In practical terms, a route that is strongly net-beneficial under climate change may still exhibit ecotoxicity burdens that either offset or dominate the Stage D credit.

E.7.4 Terrestrial Ecotoxicity (TETP)

TETP is notable because its sign reversals are concentrated in conventional recycling across all four scenarios. This makes it the clearest example of a route-specific but scenario-persistent inversion. Unlike FETP and METP, which concentrate in the emerging routes, TETP shows that even a mature closed-loop route can behave qualitatively differently once a different midpoint mechanism once a different midpoint mechanism is considered.

E.7.5 Freshwater Eutrophication, Marine Eutrophication, and Ozone Depletion (FEP, MEP, ODP)

These three indicators exhibit sign reversals concentrated in hydrolysis in all three 2050 scenarios. This pattern suggests a route-specific and prospective divergence rather than a general instability across the route set.

The interpretive implication is that hydrolysis is particularly sensitive to the non-climate burden structure of its treatment chain and receiving-system definition in future scenarios. This reinforces the main text's identification of hydrolysis as the route with the strongest receiving-system dependence.

E.7.6 Energy Resources: Fossil (FFP)

FFP is the most stable non-climate comparator in the uploaded results. It shows:

- no sign reversals relative to GWP100,
- no rank inversions relative to GWP100.

This does not mean that FFP is interchangeable with GWP100. Rather, in this case study it appears to track the same broad comparative signal more consistently than the other non-climate midpoint indicators.

E.8 Summary

Taken together, the results in this appendix support five conclusions:

1. GWP100 remains a useful primary indicator, but not a sufficient one. It captures an important part of the comparative signal, especially for landfill, reuse, and, to a lesser extent, conventional recycling.
2. The degree of qualitative stability is route dependent. Landfill and reuse are the most stable routes under cross-indicator screening. Hydrolysis is the least stable, followed by MS-FSC.

3. Sign agreement and ranking agreement should not be conflated. Many midpoint indicators preserve the net sign observed under GWP100 while still changing the comparative ordering among routes.
4. The most important non-climate divergences are concentrated in:
    a. ionising radiation,
    b. freshwater and marine ecotoxicity,
    c. terrestrial ecotoxicity,
    d. and selected hydrolysis-specific divergences in FEP, MEP, and ODP.
5. The cross-indicator screening reinforces the main text's core interpretive caution: Stage D-dominant climate benefits should not be read as a universal environmental preference. Where route comparisons depend strongly on a substitution claim, the environmental signal can change once different impact mechanisms are considered.

Appendix F — Uncertainty Propagation and Sensitivity Evidence

F.1 Purpose and Reader Guidance

This appendix provides the complete uncertainty and sensitivity evidence supporting Step 5 of the framework. It is written to be auditable as a standalone appendix: all methodological settings needed to interpret the results are stated here, and every figure/table referenced in the text is either included or marked with an explicit insertion placeholder. All plots in this appendix use p5–p95 whiskers for consistency with the sign-robustness screen. More extreme tail statistics (p2.5–p97.5) are retained in the run-summary exports and cited selectively where needed to qualify near-boundary interpretations.

Key reporting objects in this appendix are:

- Net uncertainty distributions by route and scenario (GWP100).
- Decomposed uncertainty (C3–C4 vs Stage D vs net).
- Dominance structure and sign-robustness classification using p5–p95.
- Convergence evidence (qRMSE_rel stopping criterion).
- Sensitivity ("driver") evidence identifying high-leverage foreground levers.
- Complete uncertainty input tables used for sampling (hydrolysis, MS-FSC).

Two distinct but related workflows are reported in this appendix. The first is uncertainty propagation, which asks how route outcomes are distributed when uncertain inputs are propagated through repeated Brightway LCI/LCIA evaluation. The second is sensitivity screening, which asks which sampled foreground levers are most strongly associated with output variation under that uncertainty design. The former produces route distributions, percentile bands, and convergence evidence; the latter produces driver rankings and lever–response plots.

F.2 Background Uncertainty Taxonomy (Contemporary)

Exchange-level background uncertainty was propagated for the contemporary (2025) case via Monte Carlo sampling of exchange uncertainty distributions embedded in the background inventory, while holding the contemporary foreground implementation and credit-governance ledger fixed. Because the contemporary case represents a single time-slice, there is no SSP pathway variance, and no emerging-technology foreground parameter uncertainty is propagated in this setting. The resulting Monte Carlo distributions therefore isolate within-case variability attributable to background exchange uncertainty (e.g., pedigree-style uncertainty on LCI exchanges), consistent with established LCA uncertainty practice in which parameter uncertainty is treated as a first-class contributor to dispersion and communicated using distributional summaries (medians with percentile bands) rather than single-point standard deviations.

The purpose of this subsection is to characterize the magnitude and shape of within-case variability induced by background exchange uncertainty around the contemporary reference

configuration and determine whether the qualitative conclusions for each route are sign-robust or sign-uncertain under this uncertainty layer. To support interpretation, uncertainty is reported for the net outcome (Figure F-1). All contemporary uncertainty distributions reported in this subsection reflect converged Monte Carlo runs under the qRMSE stopping criterion, as specified in the "Monte Carlo Conditioning, Convergence, and Sample-Size Adequacy (RMSE Criterion)" subsection within this appendix. In the contemporary case, uncertainty propagation is therefore a background-exchange Monte Carlo exercise around a fixed implemented foreground, not a mixed background-plus-foreground uncertainty analysis.

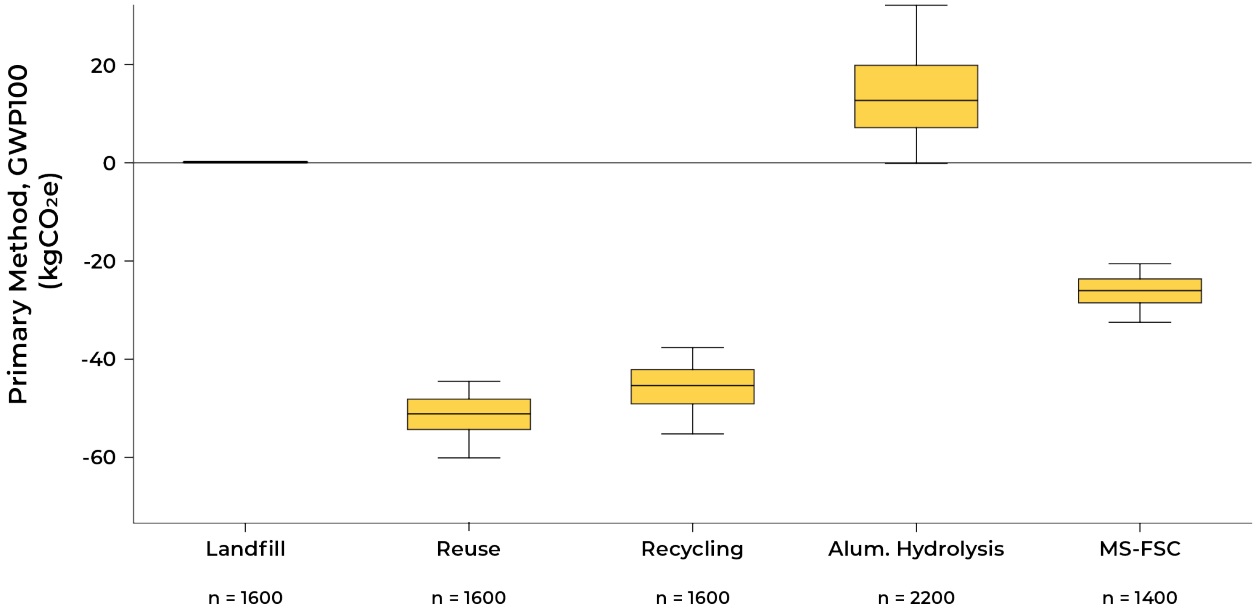


Figure F-1*: Contemporary background exchange uncertainty: net robustness across routes (GWP100)*

F.2.1 Contemporary Results: Sign Robustness and Dominance Structure

Figure F-1 provides the primary robustness screen for the contemporary (2025) case by showing how background exchange uncertainty propagates into net outcomes across routes. Table F-1 consolidates the net and decomposed (C3–C4 and Stage D) distributional results and identifies whether each route remains sign-robust under the p5–p95 interval.

The uncertainty propagation for the contemporary scenario yields the following takeaways:

- ***Landfill (baseline, no Stage D):*** Net results remain strictly positive and tightly bounded, serving as a stable reference for treatment-chain burdens in the absence of crediting.
- ***Reuse (within-sector displacement):*** Net results remain robustly net-beneficial, and the dominance ratio indicates an extreme Stage D-dominated regime. Net variability is therefore interpretable primarily as variability in the substitution credit rather than in treatment burdens.

- ***Conventional recycling (within-sector displacement):*** Net results remain robustly net-beneficial and Stage D-dominated, but with a smaller dominance ratio than reuse. The treatment burden term is non-negligible in absolute magnitude, which is relevant for interpretation beyond climate change where displaced systems can behave differently across indicators.
- ***MS-FSC (emerging route):*** Net results remain robustly net-beneficial but occupy a mixed dominance regime. MS-FSC exhibits mixed dominance because Stage D determines the sign of the net result, while C3–C4 remains large enough to materially influence the net magnitude. Therefore, both substitution-claim assumptions and treatment-chain burdens are decision-relevant for interpreting MS-FSC results.
- ***Hydrolysis (cross-sector displacement):*** Hydrolysis is the only contemporary route that is not sign-robust under p5–p95, with the net distribution straddling zero meaning the contemporary case cannot support a claim of hydrolysis being beneficial or burdensome. Hydrolysis reflects the smallest dominance ratio, indicating a burden-dominant regime in 2025, and prompting interpretation to focus on the stability of the burdens as opposed to the validity of the Stage D claim.

***Table F-1**: Contemporary (2025) background exchange uncertainty (GWP100, kg CO_2-eq per FU).*

Route	n	C3–C4 p50	C3–C4 p5	C3–C4 p95	Stage D p50	Stage D p5	Stage D p95	Net p50	Net p5	Net p95	Sign Robust? (p5–p95)	|D|/|C3–C4| (p50)	Dominance Structure
Landfill	1600	0.109	0.082	0.149	0.000	0.000	0.000	0.109	0.082	0.149	Yes	—	No Stage D (burden-only baseline)
Reuse	1600	0.012	0.008	0.019	-51.152	-60.107	-44.509	-51.140	-60.096	-44.505	Yes	4.09e+03	Stage D dominated (extreme)
Recycling	1600	2.250	1.654	3.001	-47.613	-57.466	-39.821	-45.386	-55.222	-37.642	Yes	21.161	Stage D dominated
Aluminium Hydrolysis	2200	23.536	14.790	34.717	-10.884	-20.725	-3.740	12.699	-0.150	32.103	No	0.462	Burden-dominant
MS-FSC	1400	6.366	4.932	7.988	-32.397	-38.808	-27.261	-26.035	-32.484	-20.589	Yes	5.089	Stage D dominated

F.3 Foreground Parameter Uncertainty (Prospective Emerging Technologies)

Foreground parameter uncertainty was propagated for the emerging technology routes within each SSP pathway (Figure F-2), producing pathway-specific distributions for C3–C4 burdens, Stage D effects, and net outcomes. Results are reported using median and percentile bands and interpreted through two outcome-focused questions: (i) whether the net result is sign-robust under p5–p95, and (ii) whether the route operates in a burden-dominant, mixed, or Stage D-dominant regime at central tendency. For the executed prospective analysis, this foreground-only layer was the defensible retained uncertainty layer because transformed IAM-coupled background exchange uncertainty was not propagated directly. Accordingly, the 2050 results represent within-pathway foreground variability around fixed SSP-conditioned backgrounds, rather than full joint within-scenario variance propagation.

At the input stage, the prospective foreground uncertainty runs use independent marginal sampling over the retained primitive levers. However, these draws are not interpreted as independent outputs. Each sampled realization is propagated through the shared route model and a full Brightway LCI/LCIA calculation. In hydrolysis, the sampled primitive levers $(f_{Al}, X_{Al}, L, f_{makeup}, Y_{prep}, R_{PSA}, E_{aux}, E_{therm})$ are converted into derived process quantities such as reacted aluminium, crude and usable hydrogen, aluminium hydroxide, stoichiometric water demand, electrolyte make-up and purge, and electricity demand before exchange updates are written. This preserves the internal algebraic and stoichiometric relationships encoded in the hydrolysis process model at each iteration. By contrast, MS-FSC retains a simpler structural formulation in which $pass_share$ governs the realized Stage D credit passage and $f_{transition}$ governs retained transition/overhead burden.

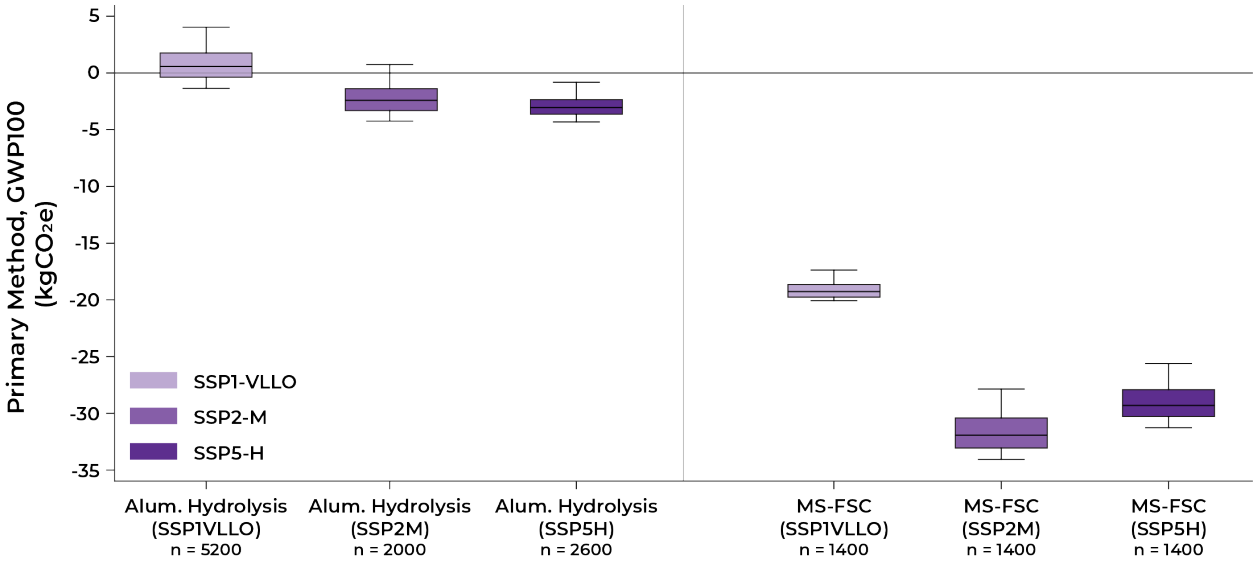


Figure F-2*: Prospective foreground-only uncertainty: net GWP100 distributions across routes and SSP pathways (p5–p95).*

F.3.1 Hydrolysis (Foreground-only Uncertainty − 2050 SSP Pathways)

Hydrolysis foreground-only uncertainty outcomes for SSP1-VLLO, SSP2-M, and SSP5-H exhibit two qualitative patterns. First, pathway choice shifts hydrolysis from sign-uncertain configurations (SSP1-VLLO and SSP2-M) to a sign-robust net benefit (SSP5-H). Second, pathway sensitivity is expressed primarily as a vertical repositioning of the net distribution relative to zero, rather than as a reversal in component directions: across SSPs, median C3–C4 remains positive and median Stage D remains negative, but changes in relative magnitude move the net distribution between sign-uncertain and sign-robust regimes. Table F-2 summarizes hydrolysis foreground-only uncertainty outcomes that underpin this interpretation.

Across the prospective pathways, hydrolysis retains the same directional structure in its component terms: median C3–C4 remains positive while median Stage D remains negative. Pathway choice nevertheless produces materially different net outcomes and robustness, indicating that the pathway-conditioned receiving-system environment changes the balance between burden-side and credit-side terms.

Under SSP1-VLLO, the net distribution is centred close to zero (net p50 ≈ +0.58) and the p5–p95 interval straddles the sign boundary, so the net outcome is not sign-robust under foreground uncertainty. Component medians indicate a narrow balance in which positive burdens (C3–C4 ≈ +4.05) are only partially offset by Stage D (≈ −3.48), consistent with a dominance ratio below unity (|SD|/|C3| ≈ 0.86). This pathway is therefore classified as a near-cancellation, sign-uncertain configuration.

Under SSP2-M, the net median shifts negative (net p50 ≈ −2.41), indicating a stronger central tendency toward net benefit than SSP1-VLLO. However, the p5–p95 interval still crosses zero (net p95 ≈ +0.73), so sign robustness is not achieved at the same reporting interval. Relative to SSP1-VLLO, the median Stage D magnitude is larger (≈ −6.01) while C3–C4 is modestly smaller (≈ +3.60), increasing separation from zero without eliminating overlap. This pathway is therefore described as credit-leaning but sign-uncertain.

In contrast, SSP5-H yields a sign-robust net benefit under p5–p95 (net p95 ≈ −0.82. Component medians show the largest separation between Stage D and C3–C4 among the pathways (Stage D p50 ≈ −6.16 vs C3–C4 p50 ≈ +3.06; |SD|/|C3| ≈ 2.02), producing a consistently negative net distribution over the reported interval. This pathway is therefore classified as credit-dominant and sign-robust.

F.3.2 MS-FSC (Foreground-only Uncertainty − 2050 SSP Pathways)

Across SSP1-VLLO, SSP2-M, and SSP5-H, MS-FSC remains sign-robust net-beneficial under the p5–p95 interval (Table F-3), indicating that the pathway-conditional net benefit persists under the modelled scale-up uncertainty (transition-overhead retention and billet eligibility/pass share). Pathway sensitivity is expressed primarily as a shift in net magnitude, consistent with the

avoided-product (Stage D) term being pathway-dependent while the same foreground uncertainty levers are applied in each case.

Across pathways, the dominance ratios confirm that Stage D governs the qualitative direction of MS-FSC results. Under SSP1-VLLO, the dominance ratio is extreme (|SD|/|C3–C4| ≈ 36), implying that the net distribution is governed almost entirely by the displacement term; the treatment-chain burden is comparatively small at central tendency (C3–C4 p50 ≈ +0.55 kg CO_2-eq, Stage D p50 ≈ −19.85 kg CO_2-eq, net p50 ≈ −19.27 kg CO_2-eq). Under SSP2-M and SSP5-H, the net benefit is larger in magnitude at the median (net p50 ≈ −31.93 and −29.31, respectively), and the dominance ratios remain well above unity (~5–6), indicating that Stage D still anchors sign while C3–C4 becomes decision-relevant for magnitude (C3–C4 p50 ≈ +6.91 in SSP2-M and +5.98 in SSP5-H). This shift is consistent with pathway-conditioned differences in (i) the receiving-system intensity of the avoided primary aluminium ingot and (ii) the upstream energy and materials burdens embedded in the MS-FSC treatment chain under each IAM-coupled inventory.

Dispersion in net outcomes is dominated by the credit-side uncertainty in SSP1-VLLO (Stage D p5–p95 width ≈ 3.82 vs C3–C4 width ≈ 0.23), while SSP2-M and SSP5-H exhibit a more balanced contribution of both uncertainty levers (Stage D widths ≈ 5.21 and 4.78; C3–C4 widths ≈ 3.11 and 2.71). This pattern is consistent with the MS-FSC uncertainty design where the efficiency of the share passed on (*pass_share*) directly scales the Stage D term (credit side), while the transition energy (*f_transition*) primarily impacts process electricity demand via transition-overhead retention (burden side). Even where burden-side dispersion becomes larger (SSP2-M/SSP5-H), the net distribution remains clearly separated from zero under p5–p95, supporting pathway-conditional sign robustness.

***Table F-2**: Hydrolysis (2050) foreground-only uncertainty by SSP pathway (GWP100, kg CO_2-eq per FU). Iteration counts (n) reflect convergence-stopped Monte Carlo budgets; convergence diagnostics are reported in §4.7.4.*

SSP Pathway	***n***	***C3–C4 p50***	***C3–C4 p5***	***C3–C4 p95***	***Stage D p50***	***Stage D p5***	***Stage D p95***	***Net p50***	***Net p5***	***Net p95***	***Sign robust? (p5–p95)***	***|D|/|C3–C4| (p50)***	***Dominance structure***
SSP1-VLLO	*5200*	*4.046*	*2.639*	*6.221*	*-3.478*	*-4.954*	*-2.211*	*0.579*	*-1.362*	*4.041*	*No*	*0.860*	*Near-cancellation (sign-uncertain)*
SSP2-M	*2000*	*3.599*	*2.740*	*4.455*	*-6.010*	*-7.033*	*-5.113*	*-2.412*	*-4.239*	*0.733*	*No*	*1.669*	*Credit-leaning (sign-uncertain)*
SSP5-H	*2600*	*3.056*	*2.016*	*4.101*	*-6.164*	*-7.026*	*-5.316*	*-3.044*	*-4.315*	*-0.821*	*Yes*	*2.016*	*Credit-dominant (sign-robust)*

***Table F-3**: MS-FSC (2050) foreground-only uncertainty by SSP pathway (GWP100, kg CO_2-eq per FU). Iteration counts (n) reflect convergence-stopped Monte Carlo budgets; convergence diagnostics are reported in §4.7.4.*

SSP Pathway	***n***	***C3–C4 p50***	***C3–C4 p5***	***C3–C4 p95***	***Stage D p50***	***Stage D p5***	***Stage D p95***	***Net p50***	***Net p5***	***Net p95***	***Sign robust? (p5–p95)***	***|D|/|C3–C4| (p50)***	***Dominance structure***
SSP1-VLLO	*1400*	*0.549*	*0.455*	*0.681*	*-19.829*	*-20.920*	*-17.097*	*-19.274*	*-21.294*	*-17.250*	*Yes*	*36.101*	*Stage D dominated*
SSP2-M	*1400*	*6.914*	*5.430*	*8.537*	*-39.329*	*-41.447*	*-36.240*	*-31.933*	*-34.859*	*-29.017*	*Yes*	*5.687*	*Stage D dominated*
SSP5-H	*1400*	*5.982*	*4.635*	*7.343*	*-35.530*	*-37.548*	*-32.767*	*-29.312*	*-31.664*	*-26.763*	*Yes*	*5.940*	*Stage D dominated*

F.3.3 Evidence-Based Range Elicitation

Table F-4 details the ranges and distributions implemented for the foreground uncertainty propagation of aluminium hydrolysis, aligned with Table 4-5 Step 4. Similarly, Table F-5 provides the values that apply to the prospective uncertainty propagation of the foreground parameters in regard to MS-FSC.

Table F-4: *Foreground Uncertainty Inputs of Prospective Hydrolysis*

Parameter	***Meaning***	***Central Value***	***Distribution***	***Bounds / Quantiles used in sampling***	***Rationale***
f_Al	Metallic Al fraction in prepared scrap	1.00	Beta-PERT (λ=4)	min 0.85, mode 1.00, max 1.00	Captures prepared-scrap heterogeneity/contamination while keeping the central case as "effectively metallic." Bounded on 0,1; downside tail represents occasional non-metallic fraction despite preparation.
X_Al	Reacted (conversion) fraction	0.95	Beta-PERT (λ=4)	min 0.85, mode 0.95, max 0.99	Represents incomplete conversion due to passivation/contact/solids-handling non-idealities. Upper bound <1 avoids implying perfect completion; mode reflects an engineered but not idealized reactor case.
f_makeup	NaOH make-up fraction (recirculation proxy)	0.20	Beta-PERT (λ=4)	min 0.05, mode 0.20, max 0.40	Encodes electrolyte retention/bleed as an operational lever without introducing speculative regeneration. Low values = high retention; high values = higher purge/poorer liquor management.
Y_prep	Scrap preparation yield into hydrolysis feed	0.85	Beta-PERT (λ=4)	min 0.70, mode 0.85, max 0.95	Captures handling rejects, sorting losses, and fines/dust removal prior to the reactor. Low tail reflects poor preparation episodes; high tail reflects high-quality preparation in mature handling conditions.
R_PSA	PSA recovery (crude → usable H_2)	0.95	Beta-PERT (λ=4)	min 0.90, mode 0.95, max 0.99	Aligns Stage D to usable hydrogen rather than crude gas. Range represents purification/conditioning losses; upper tail reflects best-case engineered operation (interpreted as "high-performance purification).

L	Working-solution intensity (L/kg Al metal basis)	150	Lognormal	GM 150, p5 80, p95 220	Strictly positive and multiplicative; represents reactor solids loading/dilution and liquor handling intensity. Lower tail = engineered higher-solids operation; upper tail = more dilute (lab-like or conservative) operation.
E_aux	Auxiliary electricity (kWh/kg prepared scrap)	0.15	Lognormal	GM 0.15, p5 0.08, p95 0.25	Aggregates mixing/pumping/solids handling/controls required at scale. Positive-only; lognormal reflects right-skew and multiplicative variability in auxiliary loads across designs/operations.
E_therm	Thermal-management electricity (kWh/kg prepared scrap)	0.05	Lognormal	GM 0.05, p5 0.02, p95 0.12	Explicitly prevents "free" exotherm control by representing cooling/heat-management auxiliaries. Positive-only; right-skew reflects variability in thermal design and operating window constraints.

NaOH molarity $C_{NaOH} = 0.240$M, liquor density $= 1.0$kg/L, stoichiometric relationships, molecular weights, and the deterministic topology of the unit-process chain (i.e., no stochastic rewiring or alternate process structure during sampling).

Table F-5*: Foreground Uncertainty Inputs of Prospective MS-FSC*

Parameter	***Meaning***	***Central Value***	***Distribution***	***Bounds / Quantiles used in sampling***	***Rationale***
f_transition	Retained share of transition/overhead energy (B term) in the C3–C4 route	0.0	Beta-PERT ($\lambda = 4$)	min 0.0, mode 0.0, max 1.0	Encodes scale-up uncertainty about how much laboratory ramp/overhead energy persists at industrial operation. Central assumes overhead removed (industrialized); uncertainty allows partial-to-full retention as a burden-side stress test.
$pass_share$	Fraction of billet output that qualifies for Stage D substitution credit	1.0	Beta-PERT ($\lambda = 4$)	min 0.7, mode 1.0, max 1.0	Implements the equivalence / eligibility gate: only qualifying billet earns Stage D credit. Central assumes full qualification; downside tail represents rejects/partial market compatibility without changing the treatment burdens.

These correspond to the Step 4 scaling logic where the prospective central case removes the transition/ramp overhead and treats billet equivalence as passing deterministically, while Step 6 re-introduces those as uncertainty levers

F.3.4 Edge Cases

This subsection documents the edge-case controls and numerical safeguards used to keep uncertainty propagation defensible in a substitution-credit model where the foreground is actively rewritten during sampling. Only a small set of independent, physically interpretable levers were sampled, and all dependent quantities were recomputed deterministically each draw to preserve closure and avoid incoherent combinations.

For hydrolysis, each draw updates throughput *(Y_prep)*, reacted aluminium *(f_Al, X_Al)*, stoichiometric water and coproduct yields (derived from reacted aluminium), usable hydrogen realization *(R_PSA)*, electrolyte make-up and purge volumes *(L, f_makeup)*, and aggregated electricity *(E_aux, E_therm)*. Infeasible draws were rejected and resampled via explicit bounds (fractions constrained to [0, 1]; lognormal intensities strictly positive). For MS-FSC, feasibility reduces to bounding (*f_transition*) and (*pass_share*) within [0, 1], since the remaining structure is deterministic in foreground-only runs.

A second edge case is “credit leakage” from embedded negative technosphere exchanges inherited by cloned background proxies inside the C3–C4 chain, which can unintentionally introduce ungoverned credits. To preserve the accounting intent that Stage D is the only deliberate credit mechanism, the prospective builders optionally strip negative technosphere “market/market group” outputs that are not waste-like or scrap-like (electricity markets ignored), and the uncertainty runners perform an FG-only traversal to audit remaining negative technosphere exchanges, optionally failing fast if any persist. Lastly, near-zero net (“offsetting-terms”) configurations are treated as an edge case because they increase sample-size sensitivity under relative convergence diagnostics; interpretation in such cases should emphasize convergence-validated tail reporting and sign-probability metrics, with detailed treatment deferred to the convergence/results subsection.

F.4 Decomposition Logic

Despite explicit uncertainty propagation and convergence validation for key runs, the Step 5 assessment does not represent the full uncertainty space implied by a prospective consequential LCA with substitution-based Stage D credits. The distributions reported in §F.3 should therefore be interpreted as conditional uncertainty under the implemented claim governance and sampled parameter set, rather than as a complete representation of epistemic uncertainty in the broader system. This section clarifies, at the level of implementation, which uncertainty layers were executed, which attribution goals could not be met, and how these limitations constrain interpretation of the prospective (2050) findings.

F.4.1 Uncertainty Coverage Achieved

In implementation, exchange-level background uncertainty propagation was executed for the contemporary (2025) inventory, where exchange uncertainty fields were retained and interpretable. In contrast, for IAM-coupled prospective databases, background exchange uncertainty was not propagated because uncertainty metadata were not consistently preserved or interpretable after IAM-driven transformations and rescaling. In scenario-coupling pipelines, rescaling routines often remove or invalidate uncertainty parameterizations, and more broadly, a rescaled future exchange does not imply that the original uncertainty distribution has been defensibly transformed to match future conditions. Under these conditions, sampling the original (or partially retained) uncertainty fields risks introducing false precision and inconsistent interpretation (e.g., mixing future mean-shifts with legacy variance structures that no longer correspond to the transformed process).

Accordingly, prospective backgrounds were treated as fixed conditional on scenario construction, and uncertainty in 2050 is represented primarily through (i) scenario uncertainty via SSP pathway contrasts and (ii) foreground parameter uncertainty for emerging technologies. This limitation implies that the reported prospective uncertainty bands may understate within-scenario variability if background exchange uncertainty would have been material for key upstream contributors, particularly for electricity- and materials-intensive processes whose upstream supply chains were strongly modified in the IAM coupling.

As a result, Step 5 parameter uncertainty evidence is currently supported by two executed components:

1. ***Contemporary Background-Exchange Uncertainty (2025):*** Monte Carlo sampling of background exchange uncertainty distributions, holding the contemporary foreground implementation fixed.
2. ***Prospective Foreground-Only Uncertainty (2050; emerging technologies):*** Monte Carlo sampling of emerging-technology parameter distributions within a fixed SSP pathway background.

A third component that was planned for attribution, prospective joint uncertainty (background + foreground), was not implemented because prospective background exchange uncertainty could not be sampled reliably under the transformed inventories. Consequently, Step 5 does not yet support full within-pathway variance attribution between background and foreground sources in 2050.

Until prospective joint attribution becomes feasible and is executed, variance attribution in 2050 remains partial. Comparisons across pathways should therefore emphasize distributional patterns that remain interpretable under a fixed-background assumption—such as dominance regimes, near-zero (offsetting-terms) behaviour, sign probabilities, and sensitivity-ranked levers—rather

than fine-grained statements about whether uncertainty is background-dominated or foreground-dominated within a given SSP.

F.4.1 Implications for Interpretation and Reporting

These implementation limitations affect what can be claimed in the prospective results and how results should be communicated:

- Prospective within-pathway uncertainty bands are conditional on the constructed IAM-coupled background and therefore reflect foreground variability around a fixed scenario inventory, not full within-scenario LCI uncertainty propagation.
- For Stage D-dominant routes, apparent sign stability under uncertainty remains conditional on the specified substitution governance (displaced-product definitions, market boundaries, marginal response representations, and substitution/equivalence ratios), because these choices largely determine the magnitude of the credit term.
- For offsetting-terms (near-cancellation) cases, such as hydrolysis in the SSP1-VLLO aligned scenario, where net results approach zero and may cross the sign boundary, probability-style metrics (e.g., P(net<0)) and convergence-validated Monte Carlo budgets are required for defensible interpretation, since central-tendency summaries can be misleading when distributions straddle zero.
- The input sampling design is based primarily on independent marginal sampling unless relationships are imposed by explicit route equations. This is methodologically appropriate for the present structured propagation and screening design, but it should not be overstated as a fully correlated joint uncertainty model.
- Sensitivity evidence reported in this appendix is association-based screening evidence. It identifies influential levers under the chosen ranges and route structure, but it is not a first-order Sobol decomposition and not a one-factor-at-a-time causal isolation exercise.

F.5 Monte Carlo Conditioning, Convergence, and Sample-Size Adequacy (RMSE Criterion)

F.5.1 Contemporary Convergence

Because Monte Carlo conclusions can be sensitive to sample size—especially when net outcomes are residuals of opposing terms—convergence was evaluated using a quantile-based RMSE diagnostic (qRMSE) computed over a fixed set of quantiles (5–95%) and normalized by the absolute median of the net result (qRMSE_rel). A relative tolerance of 1% sustained across multiple checkpoints was used as the stopping criterion (stable hits ≥ 3). This convergence design is intended to show that the reported distribution has stabilized, rather than to justify a preselected iteration budget.

For the base-routes multi-route run, convergence was assessed per route, and early stopping used a conservative worst-route criterion: at each checkpoint, the maximum qRMSE_rel across routes was taken as the governing value, and the run stopped only once this worst-case value remained

below tolerance for three consecutive checkpoints. This avoids pooled convergence that can mask a single noisy route.

The convergence profiles show a structurally informative pattern: credit-dominant routes stabilize quickly, whereas routes with small-magnitude nets under a relative error metric and routes with broader/heavier-tailed distributions require more draws. In the contemporary evidence set:

- Base routes (reuse, recycling, landfill) converged at n = 1600 under the worst-route criterion (qRMSE_rel_max ≈ 0.00307; stable hits = 3), with landfill governing the stopping time due to strict relative normalization under small net magnitudes.
- MS-FSC (contemporary) similarly met the convergence tolerance by ~1,000 iterations and sustained stability through ~1,400 (qRMSE_rel ≈ 0.55% at n ≈ 1,000; stable hits = 3 by n ≈ 1,400).
- Hydrolysis (contemporary) exhibited materially noisier convergence behaviour: qRMSE_rel fluctuated around the 1% threshold at intermediate n and only achieved sustained stability at higher iteration counts (stable hits = 3 at n ≈ 2,200, qRMSE_rel ≈ 0.84%), consistent with a wider distribution and heavier tails.

This convergence evidence supports an interpretive claim that routes operating near a sign boundary or with broad burden distributions demand higher Monte Carlo budgets to support robust statements, while highly Stage D-dominant routes can be characterized with comparatively fewer draws.

F.5.2 Prospective Convergence

Prospective convergence was evaluated using the same quantile-vector stability diagnostic applied in the contemporary runs (qRMSE computed over a fixed quantile grid, 5–95%, and normalized as qRMSE_rel using the absolute median of the net outcome as the primary scale). In the prospective setting, convergence monitoring was applied to foreground-only uncertainty propagation within each pathway-conditioned background (i.e., parameters sampled within a fixed SSP inventory). Table F-6 summarizes the checkpoint-based convergence outcomes for the prospective foreground-only Monte Carlo runs, reporting the stopping iteration and qRMSE_rel at termination for each SSP pathway under the common tolerance rule (qRMSE_rel ≤ 1% for ≥ 3 consecutive checkpoints).

Table F-6: *Prospective FG-only convergence summary (qRMSE_rel criterion; primary method; quantile grid p5–p95; tol = 1%; stable hits ≥ 3)*

Route	***SSP pathway***	***n***	***qRMSE_rel***	***qRMSE_rel (% of |median(net)|)***	***|median(net)| (denominator)***	***First checkpoint meeting tol***

Hydrolysis	SSP1-VLLO	5200	0.009090	0.909%	0.579	2800
Hydrolysis	SSP2-M	2000	0.004221	0.422%	2.412	1600
Hydrolysis	SSP5-H	2600	0.001694	0.169%	3.044	1200
MS-FSC	SSP1-VLLO	1400	0.000967	0.0967%	30.051	1000
MS-FSC	SSP2-M	1400	0.001377	0.1377%	22.194	1000
MS-FSC	SSP5-H	1400	0.000847	0.0847%	19.595	1000

Hydrolysis

Across SSP pathways, hydrolysis foreground-only runs achieved convergence under the qRMSE_rel stopping criterion (tolerance 1%; stable hits ≥ 3), but required materially different Monte Carlo budgets depending on the pathway-conditioned scale of the net result:

- ***SSP1-VLLO:*** Converged at **n = 5200** (qRMSE_rel ≈ 0.00909). This pathway required the largest budget because the hydrolysis net outcome is centered close to zero, so the relative stability metric (normalized by the median net) becomes highly sensitive: small absolute drift in the estimated 5–95% quantiles translates into a comparatively large normalized RMSE. In practical terms, SSP1-VLLO represents a near-balance case where convergence demands increase because the net distribution sits close to the sign boundary.
- ***SSP2-M:*** Converged at **n = 2000** (qRMSE_rel ≈ 0.00422). Convergence was faster than SSP1-VLLO because the pathway-conditioned net outcome is further from zero, relaxing the relative normalization and reducing sensitivity of qRMSE_rel to incremental quantile drift. This is consistent with hydrolysis having a clearer net signal under SSP2-M at central tendency, even though the upper tail still approaches (and crosses) the sign boundary in the distributional results.
- ***SSP5-H:*** Converged at **n = 2600** (qRMSE_rel ≈ 0.00169). The final qRMSE_rel is the lowest of the three pathways, indicating strong stability once converged. The slightly higher stopping iteration count relative to SSP2-M is attributable to the consecutive-hit stopping rule: qRMSE_rel briefly hovered near the 1% threshold at intermediate checkpoints, delaying accumulation of the required stable hits even though the pathway outcome is most clearly separated from zero in the hydrolysis results. This pattern is compatible with a strong net signal under SSP5-H alongside modest short-run sensitivity of the estimated quantile vector early in the run.

MS-FSC

MS-FSC convergence was evaluated under the same quantile-vector stability diagnostic used elsewhere (qRMSE over the 5–95% quantile grid, normalized to qRMSE_rel using the absolute median of the net outcome as the primary scale with 1% tolerance across ≥ 3 stable hits). In contrast to hydrolysis, where near-zero nets GWP impacts make relative stabilization stringent, MS-FSC exhibits rapid stabilization across all three SSP pathways, consistent with a strongly signed net distribution and a comparatively low-dimensional foreground uncertainty specification.

- ***SSP1-VLLO****:* Converged at **n = 1400** with qRMSE_rel ≈ 0.000967.
- ***SSP2-M****:* Converged at **n = 1400** with qRMSE_rel ≈ 0.001377.
- ***SSP5-H****:* Converged at **n = 1400** with qRMSE_rel ≈ 0.000847.

Across pathways, qRMSE_rel values for MS-FSC at convergence are well below the 1% threshold, and the stopping iteration is governed primarily by the consecutive-hit rule (stable hits accumulating across checkpoints) rather than by difficulty crossing the tolerance boundary. Substantively, these convergence profiles are consistent with the MS-FSC dominance structure: because the net distributions are far from zero and remain sign-robust under p5–p95, the relative stability metric is not inflated by small denominators, and percentile estimates stabilize under comparatively modest Monte Carlo budgets.

Cross-Scenario Interpretation of Convergence Behaviour

Two interpretation-relevant implications follow from the combined contemporary and prospective evidence. First, under a relative convergence diagnostic, Monte Carlo budgets are not transferable across routes or scenarios as configurations with net outcomes near the sign boundary require materially larger iteration counts to stabilize percentile estimates than configurations with larger-magnitude nets. Second, the pattern observed in the contemporary runs carries through to prospective emerging-technology uncertainty: strongly signed outcomes stabilize earlier, whereas near-zero (offsetting-term) outcomes require substantially larger budgets to reach the same stability tolerance. These properties justify reporting convergence status and achieved sample sizes alongside uncertainty bands, particularly for emerging technologies whose pathway-conditioned net results can shift toward the sign boundary.

F.6 Sensitivity Screening & Hotspot Identification

Sensitivity screening was used to identify which sampled foreground levers are most strongly associated with route-level net-outcome variation within fixed SSP backgrounds. The screening is global and simultaneous rather than one-factor-at-a-time. In each iteration, a full vector of lever values is sampled, patched into the route model, propagated through repeated Brightway LCI/LCIA evaluation, and recorded as a single input–output realization.

Three association-based metrics were used to summarize influence:

- Spearman's ρ, which captures monotonic association.
- Pearson's r, which captures linear association.
- Decile effect, which compares mean outputs in the top versus bottom deciles of each lever.

The ranking is screening-oriented and sorted primarily by absolute Spearman correlation. The intent of this subsection is to report what the sensitivity runs reveal about (i) the dominant drivers of output variability and (ii) which parts of the hydrolysis and MS-FSC claims are most decision-relevant for robustness.

F.6.1 Hydrolysis: Recovery and Yield as Dominant Levers

Hydrolysis sensitivity is governed primarily by recovery and yield performance parameters, which helps explain why hydrolysis is the most sign-fragile route in near-cancellation pathway settings. The strongest monotonic responses are associated with hydrogen purification recovery (R_{PSA}) and caustic make-up fraction (f_{makeup}), with secondary contributions from scrap preparation yield (Y_{prep}), reacted fraction (X_{Al}), and metallic fraction in prepared feed (f_{Al}). Because hydrolysis sensitivity uses the same primitive-to-derived quantity mapping as the foreground uncertainty runs, the screening preserves the process structure encoded in the route: lever changes are not treated as isolated exchange perturbations, but as structured changes in reacted aluminium, usable hydrogen, coproduct generation, water demand, and electricity demand.

This structure is visually explicit in the SSP1-VLLO driver panels, where the caustic make-up fraction produces the strongest slope and visibly traverses the sign boundary within the tested range, whereas SSP2-M and SSP5-H remain buffered below zero. Hydrolysis driver evidence is therefore reported by SSP in Figure F-3, Figure F-4, and Figure F-5.

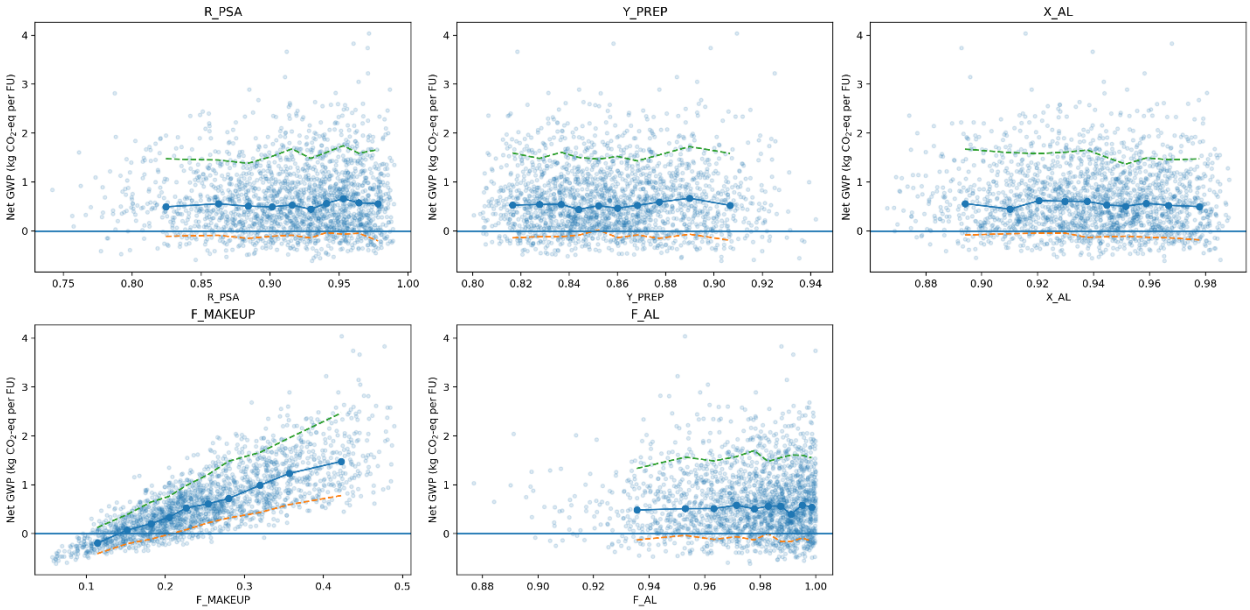


***Figure F-3**: Hydrolysis key drivers (SSP1-VLLO, 2050; gate-basis net GWP; decile medians with 10–90 bands).*

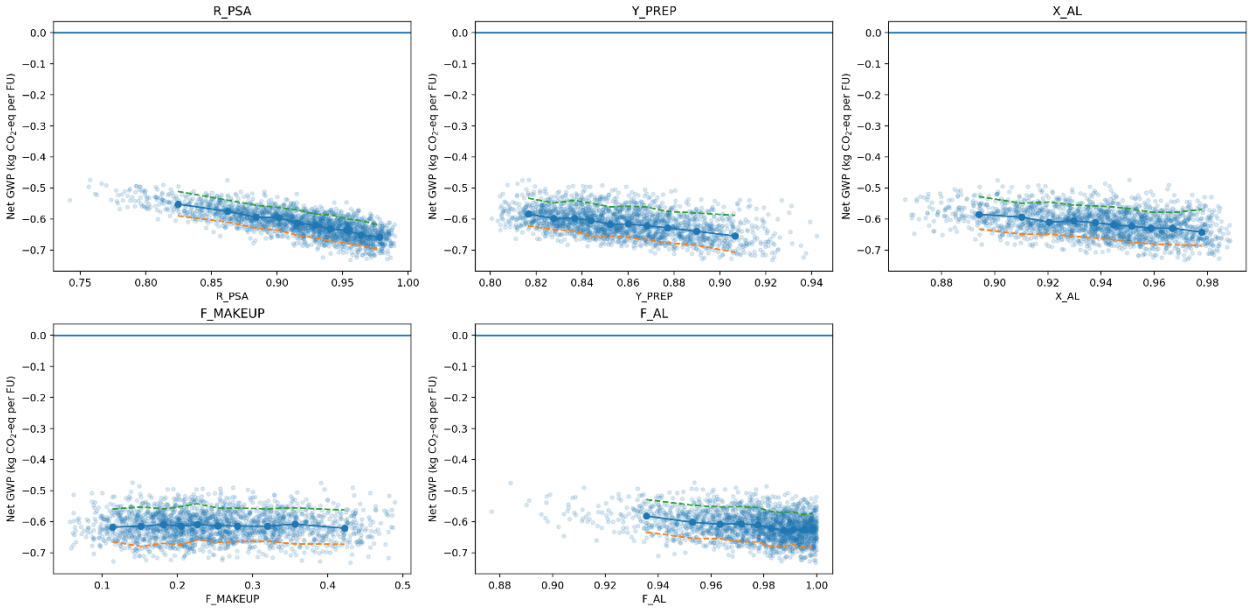


***Figure F-4**: Hydrolysis key drivers (SSP2-M, 2050; gate-basis net GWP; decile medians with 10–90 bands).*

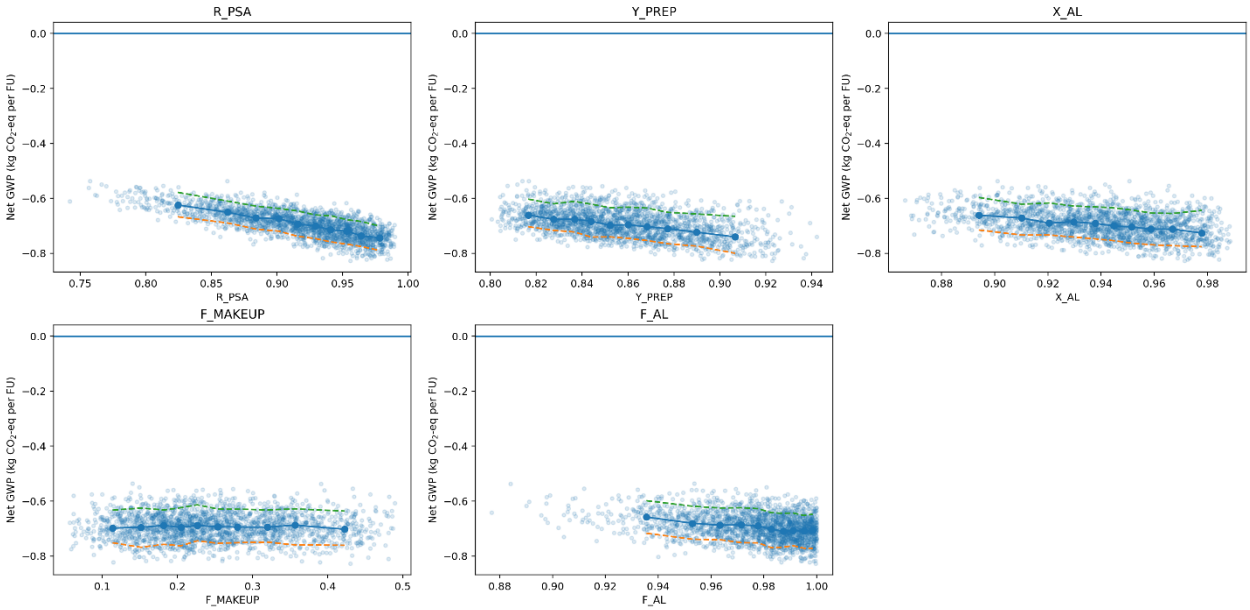


***Figure F-5**: Hydrolysis key drivers (SSP5-H, 2050; gate-basis net GWP; decile medians with 10–90 bands).*

F.6.2 MS-FSC: Substitution Strength Governs Net Response

MS-FSC sensitivity is dominated by billet qualification/eligibility share (*pass_share*), which scales the realized Stage D credit and yields an approximately linear response in net GWP within each SSP. This behaviour is consistent with the model structure: when eligibility is high, a larger fraction of output receives substitution credit; when eligibility is low, the net collapses toward the burden-only chain. The second lever, transition overhead retention (*f_transition*), produces a comparatively modest burden-side modulation by increasing electricity demand attributed to non-productive ramp/overhead energy, shifting net upward but typically not competing with billet qualification as the primary driver of outcome variance. MS-FSC driver evidence is shown by SSP in Figure F-6, Figure F-7, and Figure F-8.

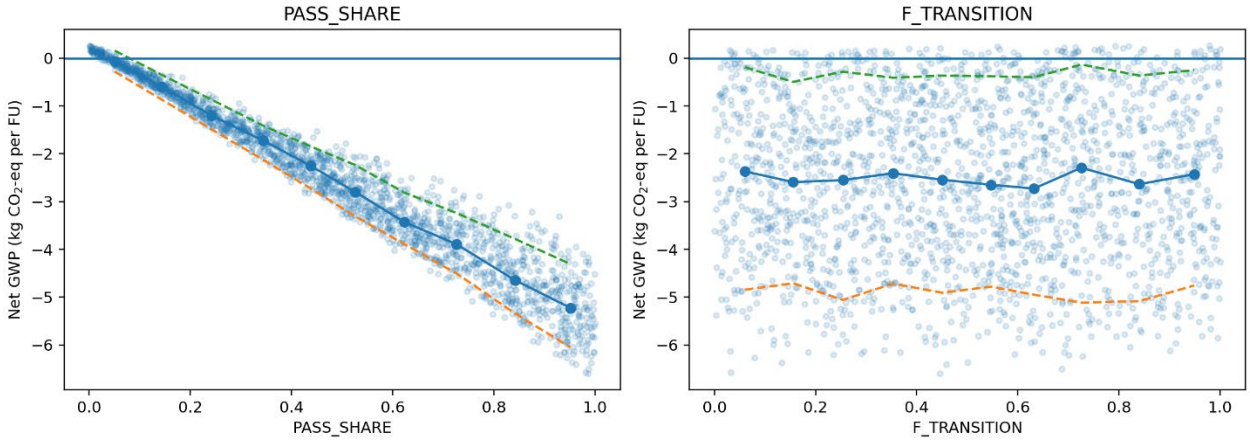


***Figure F-6**: MS-FSC key drivers (SSP1-VLLO_2050; gate-basis net GWP; decile medians with 10–90 bands).*

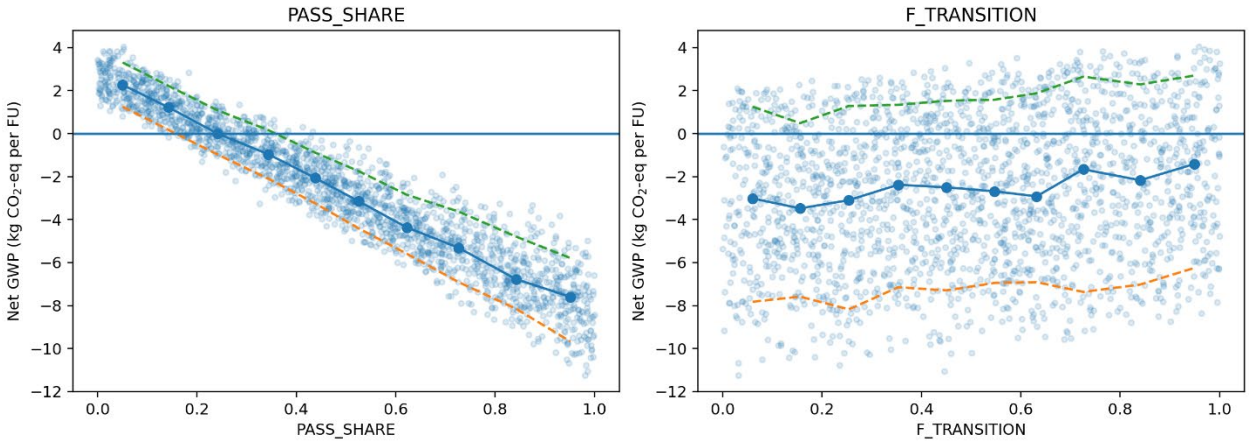


***Figure F-7:** MS-FSC key drivers (SSP2-M_2050; gate-basis net GWP; decile medians with 10–90 bands).*

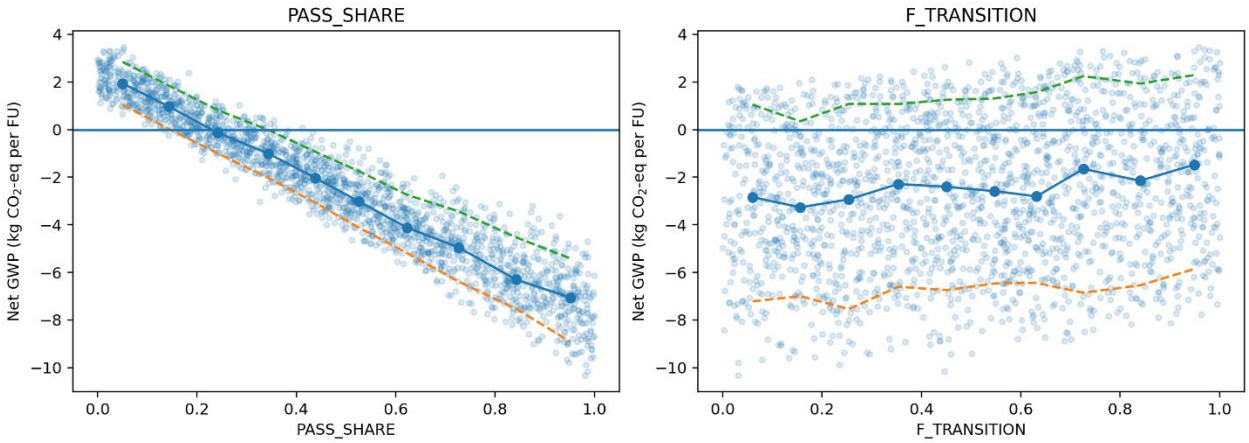


***Figure F-8**: MS-FSC key drivers (SSP5-H_2050; gate-basis net GWP; decile medians with 10–90 bands).*

The screening interpretation here is deliberately structural. MS-FSC is not represented as a full mass-balance process model in the current sensitivity design; rather, it is a route-logic model in which uncertainty enters primarily through retained overhead burden and realized credit passage. Accordingly, its driver plots should be interpreted as screening the leverage of those route-logic assumptions, not as decomposing a more detailed physicochemical process.

The Relation Between Sensitivity & Uncertainty

Taken together, the sensitivity evidence explains why uncertainty behaves differently across routes and pathways. Hydrolysis variability is governed primarily by recovery/yield performance parameters, which amplifies sign instability when pathway-conditioned net outcomes sit near zero (SSP1-VLLO and SSP2-M). MS-FSC, in contrast, is governed by substitution strength

assumptions, reinforcing that robustness statements for Stage D-governed routes are unavoidably conditional on credit governance and displacement strength even when the treatment-burden parameters are comparatively stable. Sensitivity screening does not replace the route distributions, but it helps explain which foreground assumptions most strongly condition sign robustness, magnitude, and pathway-specific fragility.

F.6.3 Scenario Uncertainty (Interpreting the SSP-overlay)

Scenario uncertainty is represented through cross-pathway contrasts under fixed IAM-coupled backgrounds. Deterministic results established that aluminium hydrolysis is receiving-system-sensitive because its dominant credited service is hydrogen rather than an in-sector aluminium product. As a result, the Stage D term inherits pathway-conditioned shifts in the displaced hydrogen system (production routes, energy inputs, and upstream supply chains). This interpretation is extended by adding within-pathway variability around each fixed SSP background through foreground parameter uncertainty propagation. Together, the SSP-overlay driver plots provide a high-resolution view of how scenario architecture changes not only central tendency, but also the decision-relevant robustness of the hydrolysis claim.

Across hydrolysis pathways, the driver overlays show that pathway sensitivity is expressed most strongly as a vertical repositioning of the net distribution relative to the sign boundary, rather than as a reversal in the direction of the component terms. In all three SSPs, the hydrolysis burden-side remains positive, and the credit-side remains negative at central tendency; however, the net buffer to zero shifts materially across SSPs. This shift is visible directly in the overlay plots: the same foreground lever ranges occupy different “distance-to-zero” regimes depending on the pathway-conditioned receiving system. In SSP1-VLLO, the hydrolysis net sits in a near-offsetting configuration where a meaningful portion of the parameter space lies close to, or above, zero. In SSP2-M, the net is centred lower, but the upper tail remains near zero. In SSP5-H, the net distribution is displaced downward enough that the same lever variation is largely absorbed without approaching the sign boundary. Scenario architecture therefore changes whether a single-sign narrative is defensible under foreground variability; it does not merely shift an average outcome.

The SSP overlays also clarify how pathway choice changes the relative importance of burden-side versus credit-side perturbations. In SSP1-VLLO, the net is sufficiently close to zero that burden-side levers become sign-relevant, meaning plausible changes in operational/scale-up parameters can move outcomes across the sign boundary. This is most clearly expressed in the hydrolysis overlays for caustic make-up fraction (*f_makeup*) and hydrogen purification recovery (*R_PSA*), which show strong within-pathway structure in SSP1-VLLO while SSP2-M and SSP5-H remain comparatively buffered (Figure F-9, Figure F-10). Interpreted through the CLCA lens, this pattern is consistent with a pathway-conditioned receiving system that supplies less net credit per unit of foreground output (i.e., a weaker displacement environment), leaving less margin for burden-side variability to be cancelled. Where the receiving system strengthens (SSP5-H in these results), the net remains decisively negative across the sampled lever ranges,

indicating that foreground uncertainty primarily tunes magnitude rather than sign (Figure F-11, Figure F-12, and Figure F-13).

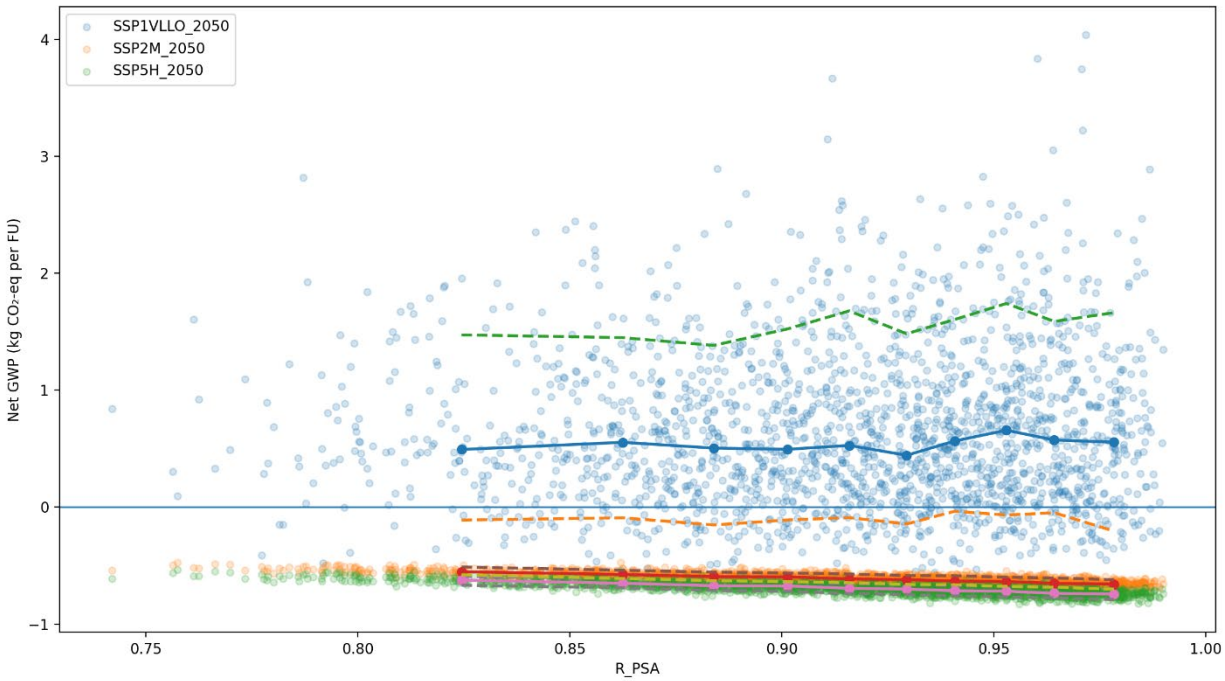


***Figure F-9**: Hydrolysis SSP overlay: net vs hydrogen purification recovery (R_PSA) (gate-basis net GWP)*

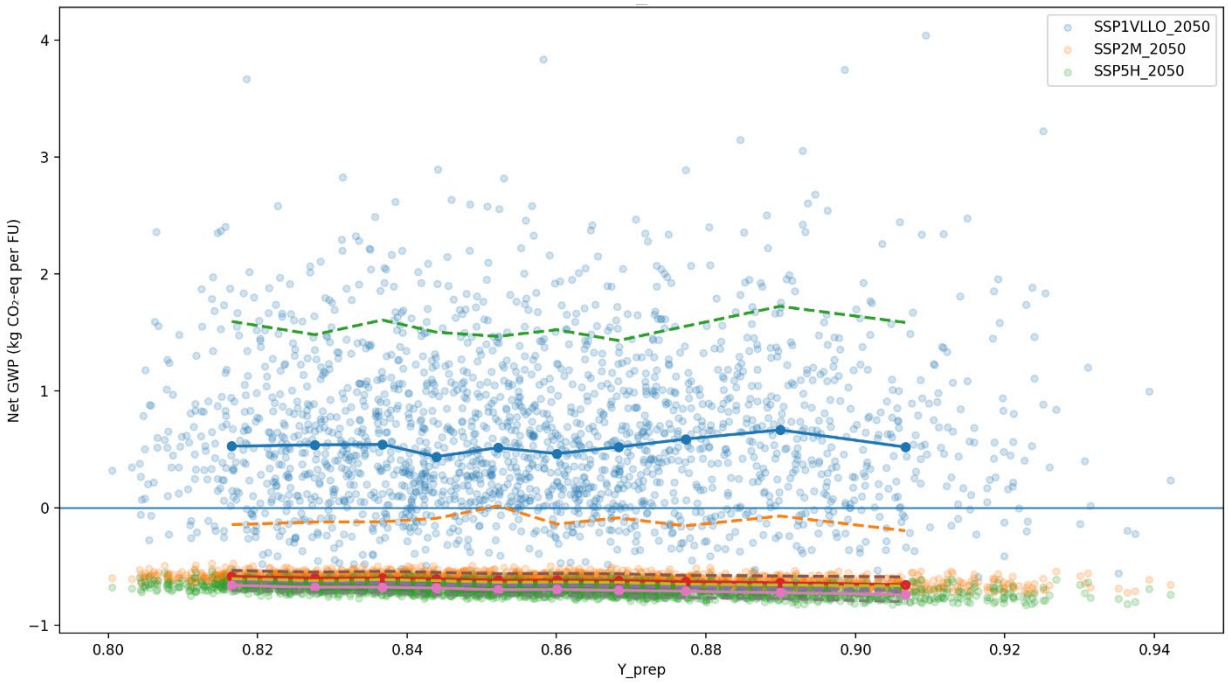


***Figure F-10**: Hydrolysis SSP overlay: net vs scrap preparation yield (Y_prep) (gate-basis net GWP)*

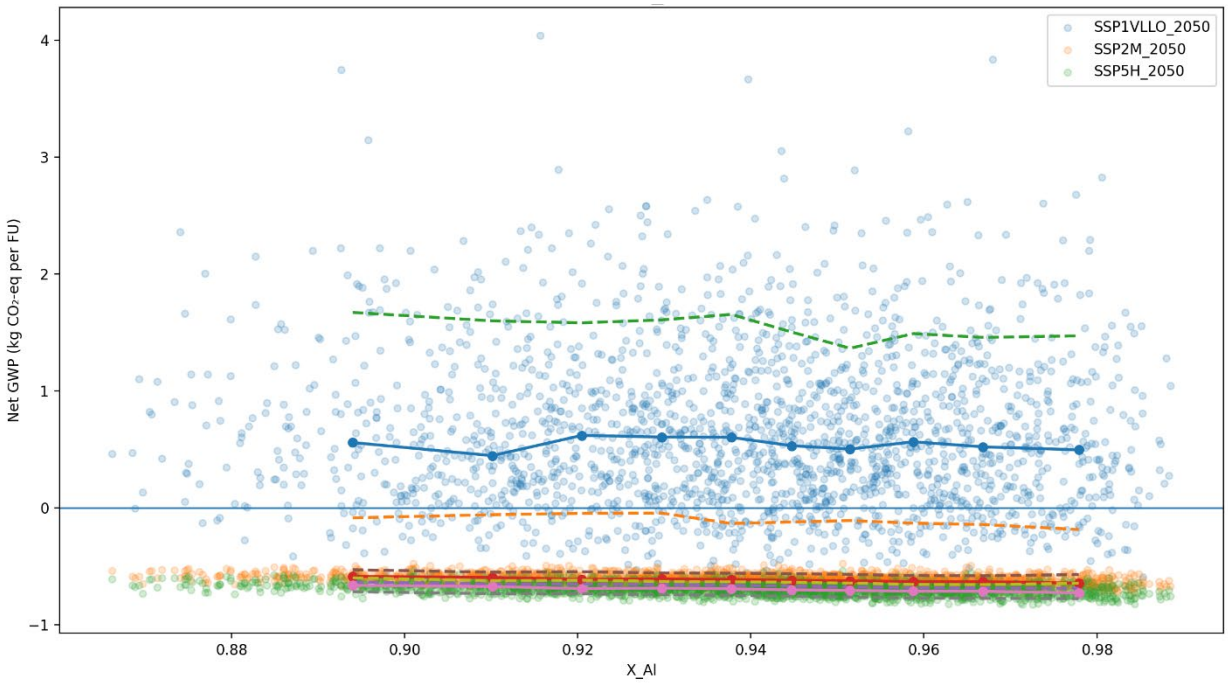


***Figure F-11**: Hydrolysis SSP overlay: net vs reacted aluminium fraction (X_Al) (gate-basis net GWP)*

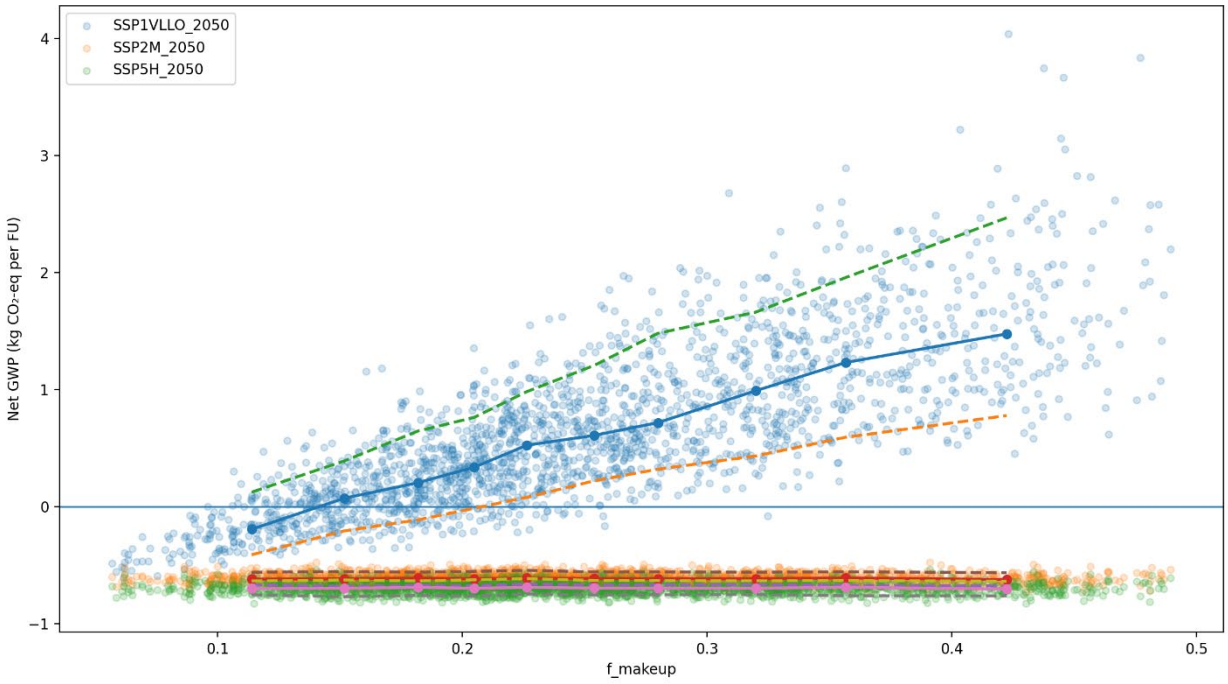


***Figure F-12**: Hydrolysis SSP overlay: net vs caustic make-up fraction (f_makeup) (gate-basis net GWP)*

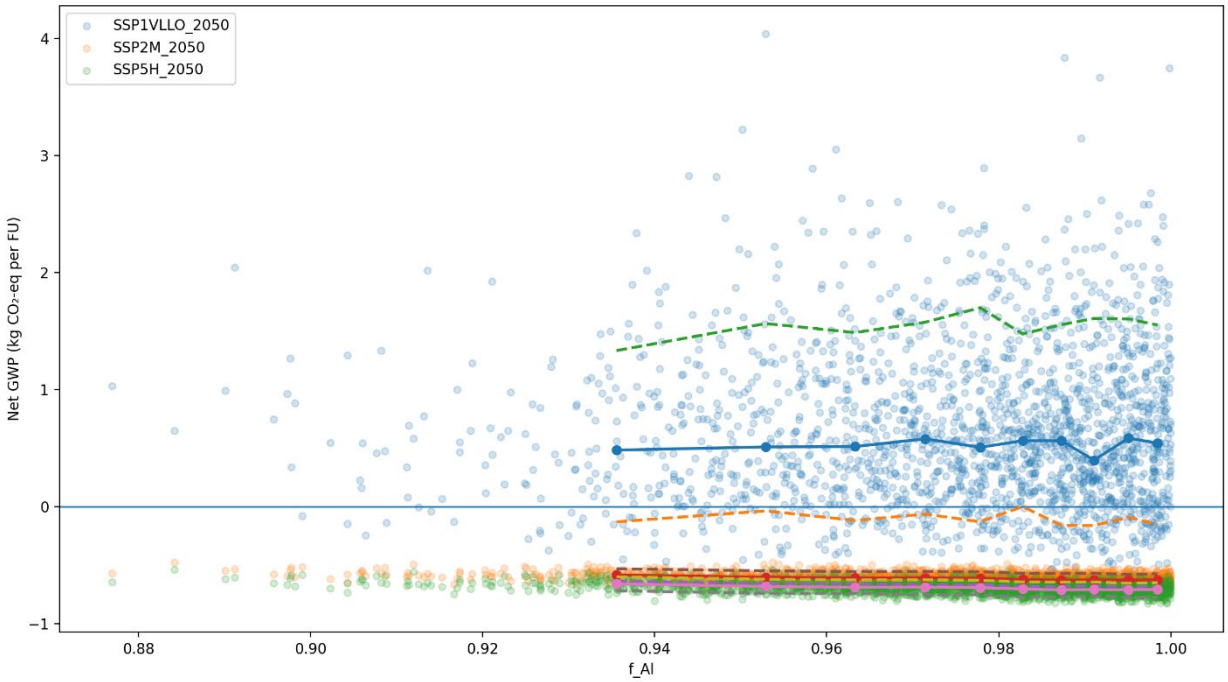


***Figure F-13**: Hydrolysis SSP overlay: net vs metallic aluminium fraction in prepared feed (f_Al) (gate-basis net GWP)*

For MS-FSC, the SSP-overlay driver plots convey a complementary scenario message: pathway choice shifts the level of net outcomes, but the structure of leverage remains dominated by a single credit-realization mechanism. Across SSPs, the qualification/eligibility share (*pass_share*) produces a near-linear net response because it scales the realized Stage D credit directly; the SSP overlays largely appear as vertical shifts around a common monotonic shape (Figure F-14). In contrast, the transition overhead retention (*f_transition*) shows a weaker, burden-side modulation of net outcomes (Figure F-15). Taken together, the SSP overlays indicate that MS-FSC is less "receiving-system fragile" than hydrolysis in sign terms under the tested conditions but remains governance-conditional in magnitude because scenario architecture still affects the implied value of the displaced product system and therefore the payoff of meeting eligibility at scale.

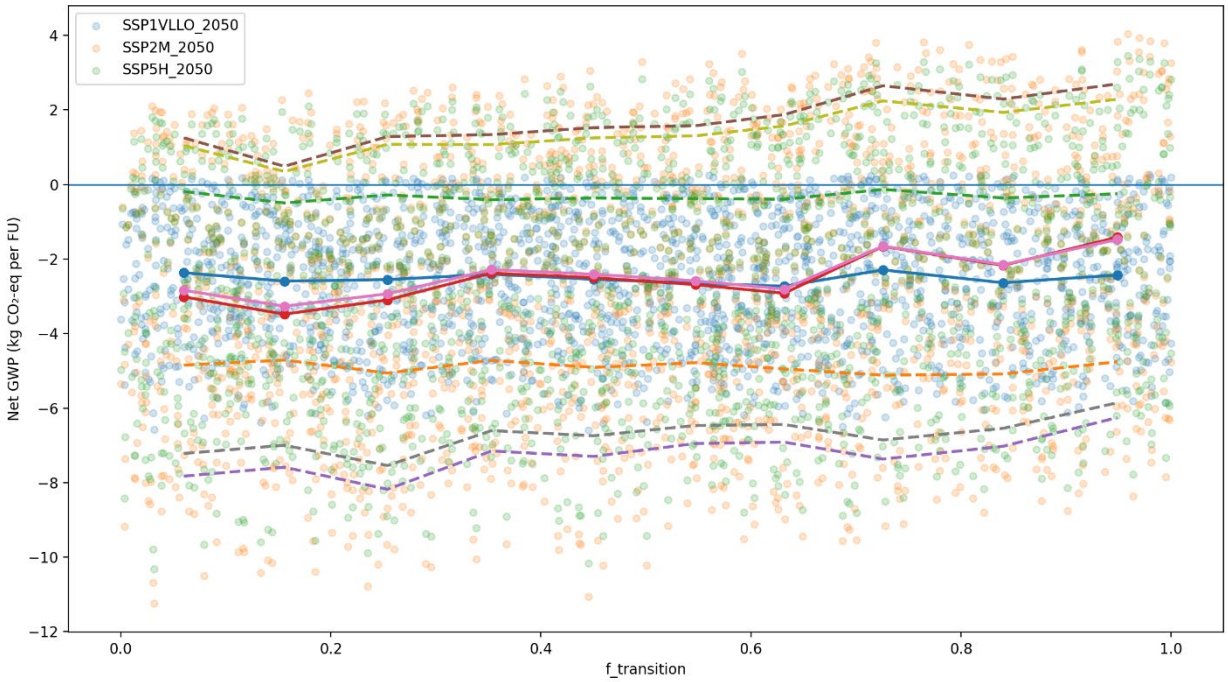


***Figure F-14**: MS-FSC SSP overlay: net vs billet qualification/eligibility share (pass_share) (gate-basis net GWP)*

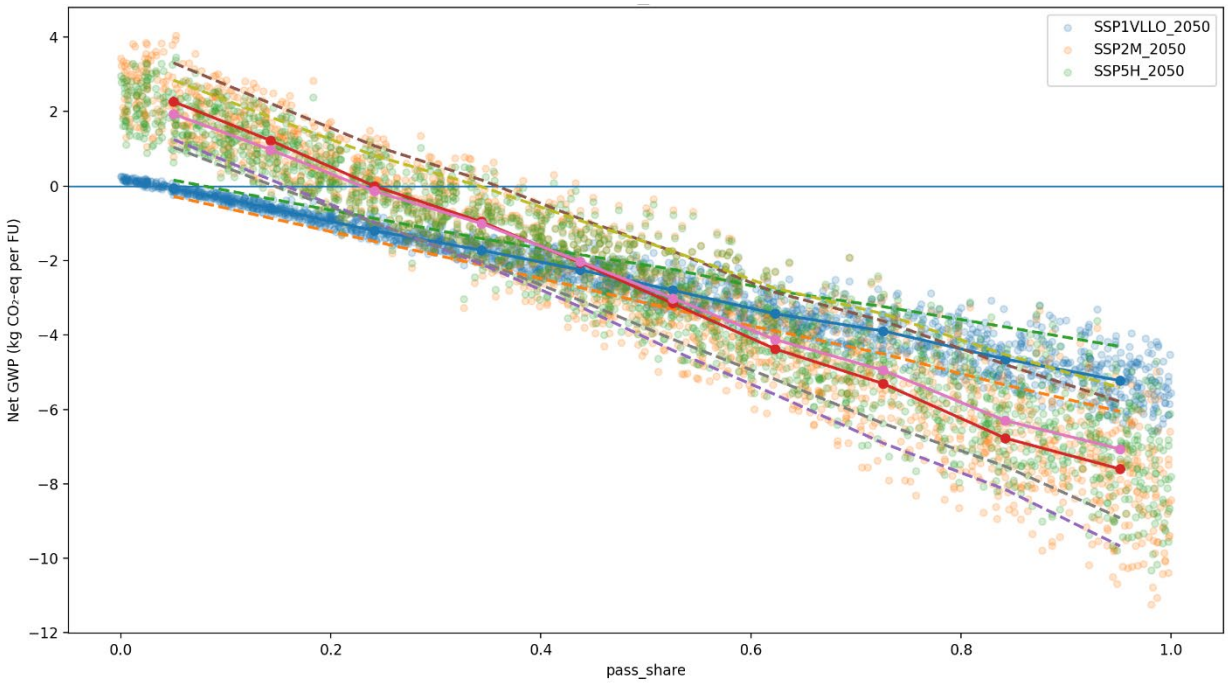


***Figure F-15**: MS-FSC SSP overlay: net vs transition overhead retention fraction (f_transition) (gate-basis net GWP)*

F.6.4 Summary

The uncertainty propagation and sensitivity analysis confirms that the interpretive logic established in the deterministic results is a structural feature of the substitution-credit framework under the implemented governance ledger. For the within-sector displacement routes, dominance persists in distribution: reuse and conventional recycling remain Stage D governed under the contemporary background-exchange uncertainty layer, meaning their net outcomes are driven primarily by the displacement claim rather than by treatment burdens. This strengthens, rather than weakens, the deterministic interpretation that rankings in credit-dominant regimes should be read as rankings of qualified substitution claims, not as direct comparisons of process emissions.

Hydrolysis emerges as the robustness hotspot across both timeframes and scenario architectures. In the contemporary (2025) case, hydrolysis occupies a burden-dominant regime and is the only route that is not sign-robust under the p5–p95 interval when background exchange uncertainty is propagated, indicating that the model cannot support a categorical "beneficial" or "burdensome" claim for hydrolysis in 2025 without additional evidence. In the prospective (2050) setting, hydrolysis remains the most fragile route in qualitative terms: under SSP1-VLLO it becomes cancellation-governed (net near zero) and remains sign-unstable under foreground parameter uncertainty, while other pathways shift the distribution further from the sign boundary. These configurations require probability-based reporting (e.g., $\mathrm{P}(\mathrm{net} < 0)$) and targeted stress tests that jointly interrogate (i) process-performance levers (yield, recovery, chemical/energy intensities) and (ii) receiving-system definitions that govern the Stage D valuation of displaced cross-sector services.

The convergence results further show that quantile-stability diagnostics are not optional, particularly for cancellation regimes. Across contemporary routes, stable percentile estimates were achieved at iteration budgets on the order of 10^3–2.2×10^3, consistent with their variance structures and distance from the sign boundary. In contrast, SSP1-VLLO hydrolysis required substantially larger Monte Carlo budgets ($\approx 5 \times 10^3$+ draws) to meet the 1% relative qRMSE criterion because the median net magnitude is small and relative stabilization becomes stringent. This implies that prospective emerging-technology uncertainty analysis must explicitly justify sample sizes and stopping rules on a route- and pathway-conditional basis rather than treating iteration budgets as transferable across configurations.

Finally, sensitivity evidence provides the mechanistic explanation for the observed uncertainty behaviour and clarifies what should be stress-tested going forward. Hydrolysis exhibits multi-lever control dominated by yield/recovery and chemical-management parameters that directly shift the balance between burdens and credits, making these levers sign-relevant whenever the pathway-conditioned net outcome approaches zero. MS-FSC, by contrast, is governed primarily by eligibility/substitution realization (e.g., pass-share), with transition/overhead retention acting as a secondary burden-side modulation. Taken together, these results motivate persistence-condition reporting, articulating the conditions under which qualitative claims hold or fail, rather

than relying solely on deterministic route rankings, particularly in Stage D-governed settings where robustness is inherently conditional on displacement governance.

Appendix G — Code Repository

https://github.com/marleydowling/marley_dowling_uw_masc_thesis