A Multi-Objective Decision-Support Framework for Early-Stage Cost–Carbon–Process-Hour Screening of Prefabricated Cold-Formed Steel Residential Framing in New Zealand

A thesis submitted in fulfilment of the

requirements for the degree

of

Master of Engineering (Engineering Thesis)

at

The University of Waikato

School of Engineering

By

Kevin Ramani

1677842

THE UNIVERSITY OF

WAIKATO

Te Whare Wānanga o Waikato

2026

Abstract

Early-stage configuration decisions in prefabricated cold-formed steel (CFS) residential framing can influence process cost, A1–A5 embodied carbon, and construction process-hours before detailed structural verification, procurement, and construction planning are complete. In the New Zealand residential construction context, these decisions are important because panel rationalisation, truss rationalisation, opening/detailing complexity, and prefabrication level affect how work is distributed across factory production, transport, logistics, and site installation. As a result, single-objective assessment is insufficient for comparing alternative framing configurations.

This thesis develops a multi-objective decision-support framework for early-stage cost–carbon–process-hour screening of prefabricated CFS residential framing in New Zealand. A project-derived residential framing package is represented using the decision vector $\mathbf{x} = [R_p, R_t, OC, P_f]$, where R_p and R_t represent panel and truss rationalisation intensity, OC represents opening/detailing complexity, and P_f represents prefabrication level. Three objectives are minimised simultaneously: CFS framing process cost, A1–A5 embodied carbon, and CFS framing process-hours. A model-calculated S2 reference configuration is used for normalisation, comparison, and robustness assessment.

An LHS-initialised NSGA-II optimisation procedure was used to generate nondominated cost–carbon–process-hour alternatives within the fixed case-study boundary. The reported optimisation run produced 8,302 feasible nondominated configurations without hard constraint violations, of which 2,695 achieved simultaneous reductions in predicted process cost, embodied carbon, and process-hours relative to S2. The remaining nondominated configurations are interpreted as trade-off alternatives rather than universal improvements over the reference case.

Model credibility was assessed through S2 reasonableness checking, coefficient-audit traceability, convergence diagnostics, sensitivity analysis, constraint testing, multi-seed stability checking, and Monte Carlo robustness assessment. The results indicate that improved performance is achieved through coordinated rationalisation, detailing simplification, and prefabrication control rather than by maximising any single variable independently. The contribution of this thesis is an applied-methodological decision-support workflow for screening prefabricated CFS residential framing configurations before detailed structural design, contractor pricing, life-cycle assessment, and construction scheduling.

Acknowledgements

I would like to express my sincere gratitude to Prof. James Lim for his guidance, patience, and critical academic support throughout this study. His feedback helped shape the direction of this study and strengthened my understanding of how engineering research must balance technical rigour, practical relevance, and methodological clarity. I am especially grateful for his encouragement to develop this work beyond a conventional optimisation exercise into a transparent decision-support framework for cold-formed steel residential framing.

I would also like to extend my sincere appreciation to Dr. Arthur Fang and Dr. Kris Roy for their valuable advice, constructive feedback, and academic insights throughout various stages of this research. Their expertise and guidance contributed significantly to the development of the methodology, interpretation of results, and overall quality of this study.

I would also like to acknowledge my fellow master's students, including Muhammad, Raj, and the wider project cohort, for creating a collaborative learning environment and for sharing ideas, challenges, and practical problem-solving approaches, particularly in relation to the use of relevant software tools. I also extend my appreciation to the PhD research community at the University of Waikato, including Humaira, Shanika, and others, for their willingness to share knowledge, research experience, and academic guidance. Special thanks are due to Mrs. Humaira Hameed for her constructive feedback, encouragement, and guidance throughout the study.

I am deeply thankful to my family for their constant support, patience, and belief in me throughout this journey. Their encouragement gave me strength during the most demanding stages of the research, particularly during the prolonged periods of modelling, analysis, writing, and revision. Their sacrifices, trust, and emotional support have been an important part of this achievement.

I would also like to thank my friends and roommates, including Smit, Tirth, Parth, and Dhruvil, for their encouragement, friendship, and support throughout this study. Their motivation and companionship helped me remain focused during the challenges of research and study preparation.

Finally, I would like to express my sincere gratitude to everyone who contributed directly or indirectly to the completion of this study. This research experience has strengthened my technical skills, critical thinking, and ability to communicate engineering decisions with transparency and confidence.

Declaration of Authorship

I, Kevin Ramani, declare that this study, entitled "A Multi-Objective Decision-Support Framework for Early-Stage Cost–Carbon–Process-Hour Screening of Prefabricated Cold-Formed Steel Residential Framing in New Zealand", is my own original work and has been completed in partial fulfilment of the requirements for the degree of Master of Engineering in Civil Engineering at the University of Waikato.

I declare that, to the best of my knowledge and belief, this study contains no material previously published or written by another person, except where proper acknowledgement and citation have been made. All sources of information, data, figures, tables, software tools, algorithms, and published works used in this study have been appropriately referenced.

I further declare that this study has not been submitted, either in whole or in part, for the award of any other degree or qualification at this or any other institution. The modelling, optimisation, analysis, interpretation, and writing presented in this study are my own work, except where assistance, supervision, or external sources have been explicitly acknowledged.

I acknowledge that this study involved the use of computational tools, including Python-based modelling, optimisation algorithms, data processing, and dashboard development. These tools were used to support the research process; however, the research design, methodological decisions, interpretation of results, and final conclusions remain my own responsibility.

Author: Mr. Kevin Ramani

Student: University of Waikato

Year: 2026

Table of Contents

Contents

List of Tables

List of Figures

List of Symbols and Abbreviations

Symbol	Definition	Unit / Type
$\boldsymbol{x}$	Candidate CFS framing configuration / design vector	Vector
$\boldsymbol{x_i}$	(i)-th candidate solution in the evaluated population or Pareto set	Vector
$\boldsymbol{x_{S2}}$	S2 model-calculated reference design vector	Vector
$\boldsymbol{\Omega}$	Feasible design domain after variable bounds and feasibility screening	Set
$\boldsymbol{F(x)}$	Raw multi-objective function vector	Vector
$\boldsymbol{C(x)}$	Raw CFS framing process cost	NZD
$\boldsymbol{E(x)}$	Raw CFS framing A1–A5 embodied carbon	$kgCO_{2e}$
$\boldsymbol{T(x)}$	Raw CFS framing process-hours	h
$\boldsymbol{R_p}$	Panel rationalisation intensity	Dimensionless
$\boldsymbol{R_t}$	Truss rationalisation intensity	Dimensionless
$\boldsymbol{OC}$	Opening/detailing complexity index	Dimensionless
$\boldsymbol{P_f}$	Prefabrication level	Fraction
$\boldsymbol{S_p}$	Derived panel scheme / panel rationalisation reporting category	Index
$\boldsymbol{S_t}$	Derived truss scheme / truss rationalisation reporting category	Index
$\boldsymbol{A_f}$	Case-study building footprint area	m^2
$\boldsymbol{A_e}$	Case-study envelope/framing area	m^2
$\boldsymbol{N_p}$	Number of wall panels in the case-study framing package	Count
$\boldsymbol{N_t}$	Number of roof trusses in the case-study framing package	Count
$\boldsymbol{W(x)}$	Modelled CFS steel mass for candidate solution (x)	kg
$\boldsymbol{W_{S2}}$	S2 reference CFS steel mass	kg
$\boldsymbol{C_{S2}}$	S2 reference CFS framing process cost	NZD
$\boldsymbol{E_{S2}}$	S2 reference CFS framing A1–A5 carbon	$kgCO_{2e}$
$\boldsymbol{T_{S2}}$	S2 reference CFS framing process-hours	h
$\boldsymbol{\overline{C}(x)}$	Normalised CFS framing process cost relative to S2	Dimensionless
$\boldsymbol{\overline{E}(x)}$	Normalised CFS framing A1–A5 carbon relative to S2	Dimensionless

$\overline{T}(x)$	Normalised CFS framing process-hours relative to S2	Dimensionless
$\Delta C(x)$	Absolute change in CFS framing process cost relative to S2	NZD
$\Delta E(x)$	Absolute change in CFS framing A1–A5 carbon relative to S2	$kgCO_{2e}$
$\Delta T(x)$	Absolute change in CFS framing process-hours relative to S2	h
ΔC	Signed percentage change in cost relative to S2	%
ΔE	Signed percentage change in carbon relative to S2	%
ΔT	Signed percentage change in process-hours relative to S2	%
$\mathbf{P}$	Pareto set / nondominated solution set	Set
A_g	Nondominated archive at generation (g)	Set
$\mathbf{g}$	NSGA-II generation index	Count
N_P	NSGA-II population size	Count
N_G	Total number of NSGA-II generations	Count
N_{seed}	Number of independent optimisation seeds used for stability testing	Count
(N_{ND})	Number of nondominated solutions	Count
(N_F)	Number of feasible solutions	Count
(N_U)	Number of unique design vectors	Count
N_{MC}	Number of Monte Carlo simulations	Count
(P_5)	5th percentile from Monte Carlo output distribution	Same as objective
(P_{50})	50th percentile / median from Monte Carlo output distribution	Same as objective
(P_{95})	95th percentile from Monte Carlo output distribution	Same as objective
(ρ)	Spearman rank correlation coefficient	Dimensionless
$p(C < S2)$	Probability that a candidate solution has lower cost than S2 under paired uncertainty sampling	Probability
$p(E < S2)$	Probability that a candidate solution has lower carbon than S2 under paired uncertainty sampling	Probability
$p(T < S2)$	Probability that a candidate solution has lower process-hours than S2 under paired uncertainty sampling	Probability
p_{better}	Probability that a candidate solution outperforms S2 under the selected robustness criterion	Probability

R_{rob}	Robustness rank assigned after Monte Carlo comparison	Rank
S_{rob}	Robustness score used for robustness-preferred solution selection, where applied	Dimensionless
w_C	Stakeholder weight assigned to cost in ranking	Fraction
w_E	Stakeholder weight assigned to embodied carbon in ranking	Fraction
w_T	Stakeholder weight assigned to process-hours in ranking	Fraction
S_i	TOPSIS ranking score for candidate solution (i), where used	Dimensionless
RC_i	Relative closeness coefficient in TOPSIS ranking, where used	Dimensionless
U_i	Utility or preference score for candidate solution (i), where used	Dimensionless
S2	Model-calculated reference CFS framing configuration used for comparison	Reference case

Abbreviations

Abbreviation	Full form
A1-A3	Product stage modules: raw material supply, transport to manufacturing, and manufacture
A4	Transport to construction site
A5	Construction and installation stage
AISI	American Iron and Steel Institute
AS/NZS	Australian/New Zealand Standards
BIM	Building Information Modelling
BRANZ	Building Research Association of New Zealand
CAD	Computer-Aided Design
CFS	Cold-Formed Steel
CO_{2e}	Carbon dioxide equivalent
CPM	Critical Path Method
CSV	Comma-Separated Values
CV	Coefficient of variation (%)
DEAP	Distributed Evolutionary Algorithms in Python
DfMA	Design for Manufacture and Assembly
DSM	Direct Strength Method
EN	European Standard
EPD	Environmental Product Declaration
GHG	Greenhouse gas
HV	Hypervolume indicator used to assess Pareto-front convergence and spread
IGD	Inverted Generational Distance used for grid-reference comparison
LCA	Life-Cycle Assessment
LHS	Latin Hypercube Sampling
MOO	Multi-Objective Optimisation
MOEA	Multi-Objective Evolutionary Algorithm
NSGA-II	Non-dominated Sorting Genetic Algorithm II
NZD	New Zealand dollars
OAT	One-at-a-Time sensitivity analysis

PRCC	Partial Rank Correlation Coefficient
SPEA2	Strength Pareto Evolutionary Algorithm 2
TOPSIS	Technique for Order Preference by Similarity to Ideal Solution
VS Code	Visual Studio Code

Chapter 1: Introduction

1.1: Opening Problem Statement

Panelised cold-formed steel (CFS) residential framing involves a set of early configuration decisions that are usually made before their combined cost, carbon, and process-hour consequences are fully visible. A CFS designer or fabricator may simplify wall-panel families, rationalise roof-truss span groups, standardise opening and detailing conditions, or increase the level of off-site prefabrication. Each decision can improve one part of the production system while increasing burden elsewhere. Higher prefabrication may reduce site process-hours but increase factory processing, logistics coordination, handling, and transport requirements. Similarly, stronger rationalisation may improve repetition and setup efficiency but can introduce oversizing, trimming, or coordination penalties.

The engineering problem addressed in this thesis is therefore not whether CFS framing is universally preferable to timber, concrete, or hot-rolled steel. The problem is narrower and more practical: once a CFS framing system has been selected, how can early framing-configuration alternatives be compared in terms of process cost, A1–A5 embodied carbon, and process-hours before detailed procurement, pricing, and construction planning are finalised? This creates a multi-objective decision problem because cost, carbon, and process-hours do not necessarily improve in the same direction.

This study responds to that problem by developing a bounded decision-support framework for a real-project-derived CFS residential framing case study. The framework evaluates alternative configurations using panel rationalisation intensity, truss rationalisation intensity, opening/detailing complexity, and prefabrication level, and compares their performance against a model-calculated S2 reference configuration (IPCC, 2023; Lui Brame, 2024; Ministry of Business, Innovation and Employment, 2020a).

The New Zealand relevance of this problem is strengthened by the continuing scale of residential construction activity. Stats NZ reported that 36,619 new dwellings were consented in Aotearoa New Zealand in the year ended December 2025, an increase of 9.0% from the year ended December 2024. This does not imply that all dwellings use CFS framing, but it shows the scale of repeated residential building activity in which framing standardisation, prefabrication, cost control, and upfront-carbon comparison are practically relevant (Stats NZ, 2026).

Panelised CFS framing is relevant to this problem because it combines structural framing decisions with production and assembly decisions. Factory-made wall panels, roof trusses, repeatable component families, and digital production workflows can improve dimensional control and reduce some site activities compared with more site-dependent framing approaches. However, these advantages cannot be assumed automatically. Higher prefabrication may reduce site process-hours while increasing factory processing, setup, logistics coordination, transport, and handling requirements. Similarly, CFS carbon performance depends on steel mass, steel emission factors, transport distance, waste assumptions, fabrication processes, and the adopted A1–A5 boundary (Grubb, 2001).

The figure uses selected extracts from the fabrication drawing package. Panel A shows the fixed S2 wall-panel layout and panel identifiers used to motivate panel rationalisation intensity R_p. Panel B shows an example wall-panel opening detail where lintels, angle plates, jamb studs, trimming, and local fit-up motivate the opening/detailing complexity variable OC. Panel C shows the roof-truss layout and truss-type diversity used to motivate truss rationalisation intensity R_t. Prefabrication level P_f is not directly visible as geometry; it represents the factory–site work split applied to the wall-panel and truss package. The figure anchors the optimisation variables in project-derived CFS fabrication information and is not presented as a structural redesign drawing.

The research challenge is that several early CFS residential framing decisions are interdependent and cannot be reliably optimised through manual judgement alone. Reducing the number of panel-width classes may improve factory repetition, but excessive rationalisation may introduce oversizing, trimming, or constructability penalties. Rationalising roof-truss families may improve standardisation, but it may not reduce cost, carbon, and process-hours equally. Opening/detailing complexity also affects lintels, jamb studs, trimmers, CNC cutting, service penetrations, and site fit-up. Prefabrication level further changes the balance between factory preparation and site installation.

The research problem is therefore formulated as a three-objective CFS framing configuration problem. The purpose is not to identify one universally best framing configuration, but to generate and interpret nondominated alternatives that expose the trade-off between CFS framing process cost, CFS framing A1–A5 embodied carbon, and CFS framing process-hours. In practice, these decisions may be guided by fabricator experience, supplier preference, spreadsheet estimating, or isolated cost and time judgement. This thesis formalises the comparison so that alternatives can be evaluated consistently against the S2 reference case and inspected for cost, carbon, process-hour, and robustness implications.

$$\min_{x \in \Omega} F(x) = [C(x), E(x), T(x)]$$

where $C(\mathbf{x})$ is CFS framing process cost in New Zealand dollars, $E(\mathbf{x})$ is CFS framing A1–A5 embodied carbon in $kgCO_2e$, $T(\mathbf{x})$ is CFS framing process-hours, and Ω is the feasible design domain.

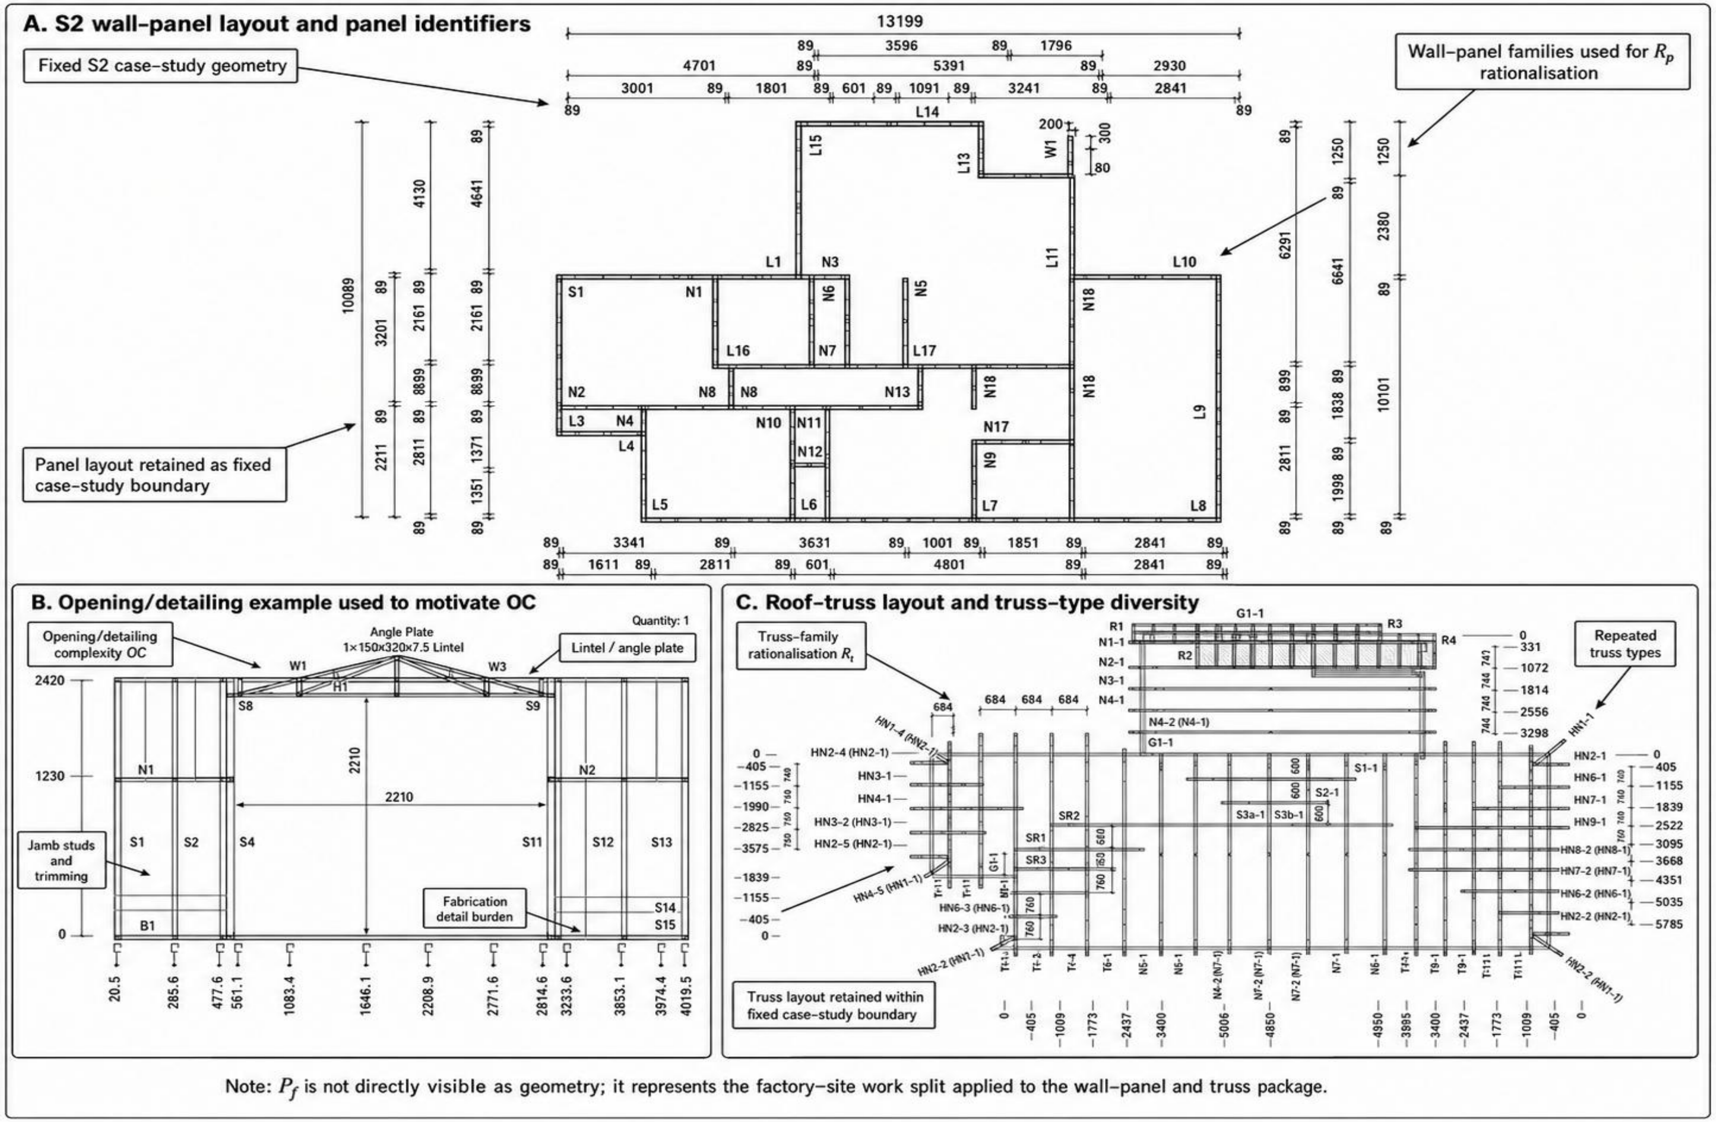

Figure 1-1 Project-derived CFS framing configuration decisions represented in the optimisation framework.

This study develops a Python-based decision-support framework for a case-study CFS residential building using Latin Hypercube Sampling (LHS)-initialised NSGA-II, followed by Monte Carlo uncertainty analysis and dashboard-based visualisation. LHS is employed for initial sampling of the design space – used to stratify the input variables used for optimisation, according to sampling logic introduced by McKay, Beckman, and Conover (McKay et al., 1979).

In this study, decision support refers to a structured early-stage comparison process for CFS designers, fabricators, project engineers, construction managers, and sustainability assessors who need to evaluate framing-configuration alternatives before detailed fabrication, procurement, and construction planning decisions are finalised. The detailed scope and modelling boundary are defined in Section 1.4.4.

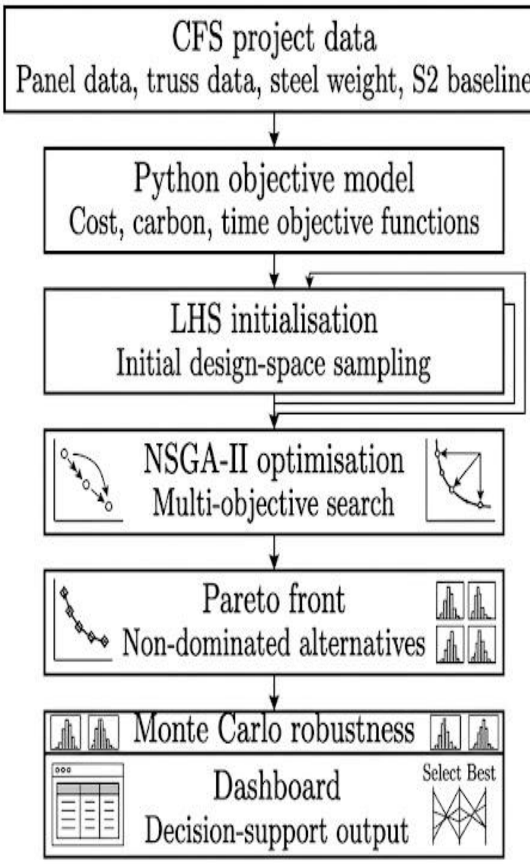

Figure 1-2 Research workflow from CFS project data to optimisation and dashboard-based decision support.

The workflow shows how project-derived CFS framing data are converted into a fixed S2 reference, evaluated through cost, A1–A5 carbon, and process-hour response functions, searched using LHS-initialised NSGA-II, tested through sensitivity and Monte Carlo robustness analysis, and translated into dashboard-based decision support.

1.2: Research Background

Cold-formed steel framing is a lightweight structural system formed from thin galvanised steel sheets into studs, tracks, joists, and truss members. In residential construction, CFS framing is commonly organised through panelised wall frames and roof-truss assemblies, which makes it strongly affected by production and assembly decisions. Digital modelling, roll-forming, cutting, punching, labelling, factory assembly, transport, and site installation create a workflow in which early configuration choices influence both factory and site performance (American Iron and Steel Institute, 2022; Schafer, 2008; Standards Australia / Standards New Zealand, 2018; Yu & LaBoube, 2010).

For this reason, the thesis focuses on four configuration variables: panel rationalisation intensity R_p, truss rationalisation intensity R_t, opening/detailing complexity OC, and prefabrication level P_f. These variables represent production and detailing decisions rather than architectural redesign or live structural member sizing. The fixed structural boundary is developed in Chapter 3, while Chapter 1 uses these variables only to define the research problem (Blismas et al., 2006; Veljkovic & Johansson, 2006).

The key background issue is that the variables do not influence the objectives independently. Increasing prefabrication can reduce site process-hours but increase factory and logistics effort. Rationalising panel or truss families can improve repetition but may create material, detailing, or coordination penalties. Simplifying opening/detailing conditions can improve manufacturability and site fit-up, while more customised detailing can increase local reinforcement, cutting, waste, and installation burden. These coupled effects justify a multi-objective response-model approach rather than a single-objective comparison (Shahzad et al., 2024).

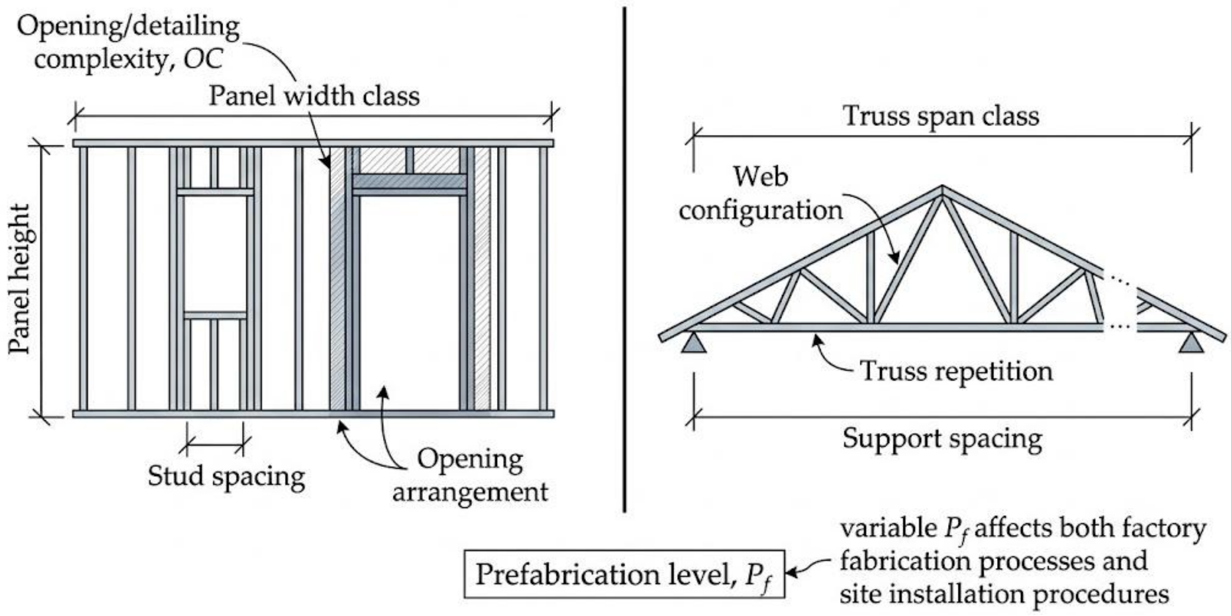

Figure 1-3 Annotated CFS panel and truss schematic with variables.

Because of these trade-offs, a multi-objective optimisation approach is necessary; The objective is to compare these objectives without imposing subjective weights. The non-dominated sorting and elitism of NSGA-II is well suited for this task as it yields a set of Pareto optimal alternatives instead of a weighted solution. The original NSGA-II paper by Deb et al is the methodological reference that is essential to the algorithm (Deb et al., 2002). NSGA-II has also been employed to address construction decision-support problems, including prefabrication cost, duration, and carbon trade-offs with respect to cost, duration and carbon emissions, further underlining its relevance in the construction domain (Wang et al., 2024). The notion of a dashboard as a translation layer is presented in this study, which places the optimisation model and the decision-maker in a dialogue. It is designed to render Pareto alternatives interpretable via trade-off plots, representative solutions, stakeholder ranking and uncertainty outputs. The dashboard is therefore positioned as an interpretation and communication layer that supports early-stage comparison of CFS framing alternatives before detailed structural verification, procurement, and project-specific construction planning are undertaken.

1.3: Research Gap

The research gap addressed in this thesis is an integration gap between CFS structural optimisation, prefabrication research, embodied-carbon assessment, construction process modelling, uncertainty analysis, and visual decision support. Existing CFS optimisation research provides important insight into member-level and section-level performance, but it does not directly address early-stage production-configuration decisions for a project-derived panelised residential CFS framing package. Conversely, prefabrication and DfMA studies discuss standardisation, production efficiency, logistics, and site productivity, but they rarely connect these mechanisms to CFS-specific panel rationalisation, truss rationalisation, opening/detailing complexity, and prefabrication level in one cost–carbon–process-hour optimisation framework (Liang et al., 2022) (Xiao et al., 2022)

Prefabrication research provides a second body of relevant literature, but much of it evaluates prefabrication through isolated performance outcomes rather than integrated optimisation. Several studies examine specific benefits such as construction waste reduction, cost effects, productivity improvement, or environmental performance. For example, Jaillon et al. (2009) examined the waste reduction potential of prefabrication in Hong Kong building construction, while Mao et al. (2016) analysed cost implications of sustainable off-site construction through multiple case studies (Aghasizadeh et al., 2022; Jaillon et al., 2009). Blismas et al. (2006) also argued that off-site production evaluated more holistically rather than only through direct cost, because benefits and constraints extend across labour, logistics, quality, coordination, and production efficiency (Blismas et al., 2006). These studies are useful for understanding prefabrication performance, but they generally do not provide a project-specific optimisation framework for residential CFS framing in which panel rationalisation, truss rationalisation, opening/detailing complexity, and prefabrication level are evaluated simultaneously. Few studies link these production variables directly to CFS framing process cost, CFS framing A1–A5 carbon, and construction process-hours in a single optimisation workflow.

Multi-objective evolutionary algorithms have been widely applied in building and construction optimisation because they can expose trade-offs between conflicting objectives instead of forcing the designer to select one fixed weighted solution. NSGA-II is one of the most established methods in this area because it uses elitist non-dominated sorting and diversity preservation to generate Pareto-optimal solution sets (Deb et al., 2002). Building optimisation reviews show that simulation-based and computational optimisation methods have been widely used for sustainable building design, energy performance, cost, and environmental objectives (Evins, 2013; Nguyen et al., 2014). More recent prefabrication optimisation research has also used NSGA-II to optimise prefabricated component combinations using cost, duration, and carbon emissions as simultaneous objectives (Wang et al., 2024). However, these applications commonly focus on broader building design variables, energy systems, architectural envelope parameters, or generic prefabricated component selection. The specific CFS framing variable set used in this study—panel rationalisation intensity, truss rationalisation intensity, opening/detailing complexity, and prefabrication level—has not been widely integrated into a single NSGA-II framework for panelised residential CFS decision-making. This creates a methodological gap between general building optimisation research and the practical configuration decisions required in CFS residential prefabrication.

A further gap exists in the treatment of uncertainty, sensitivity, and decision support. Many optimisation studies report deterministic Pareto fronts or static comparison tables, but practical construction decisions are affected by uncertain steel costs, labour productivity, transport assumptions, carbon factors, weather exposure, and factory/site coordination. Latin Hypercube Sampling provides a structured method for improving design-space coverage through stratified sampling, and it is widely used in computer experiments and uncertainty analysis following the work of McKay et al. (1979) (McKay et al., 1979; Mckay et al., 2000). Similarly, Monte Carlo analysis and sensitivity methods are important because deterministic solutions may appear optimal under one parameter set but become less attractive under uncertainty. Decision-support literature also emphasises that interactive visualisation can help users explore, compare, filter, and interpret complex data rather than passively reading static outputs; Heer and Shneiderman (2012), for example, describe visual analytics as an iterative process requiring filtering, selection, navigation, coordination, annotation, and guidance (Heer & Shneiderman, 2012). In the context of this study, this means that a Pareto front alone is not sufficient. Practitioners need a dashboard or interface that translates optimisation outputs into interpretable alternatives, such as minimum-cost, minimum-carbon, minimum-time, balanced, and robust within the tested uncertainty ranges.

The novelty of this study is applied-methodological rather than algorithmic. NSGA-II, LHS, Monte Carlo simulation, TOPSIS, and dashboard visualisation are established methods. The contribution lies in adapting and integrating these methods around a project-derived prefabricated CFS residential framing decision problem in the New Zealand context. Specifically, the thesis links four CFS-specific configuration variables, R_p, R_t, OC, and P_f, to process cost, A1–A5 embodied carbon, and process-hours through a coefficient-audited response model and then tests the resulting alternatives through S2-relative comparison, sensitivity analysis, Monte Carlo robustness assessment, stakeholder ranking, and dashboard-supported interpretation.

1.4: Aim, Objectives, and Research Questions

The following aim and research questions are structured to address this integration gap. They move from defining the bounded CFS framing decision problem, to modelling the cost–carbon–process-hour response, to generating and interpreting Pareto alternatives, to testing robustness, and finally to translating the results into practical decision support.

1.4.1: Research Aim

The aim of this research is to develop and evaluate a bounded multi-objective decision-support framework for early-stage prefabricated CFS residential framing configuration. The framework compares alternative configurations in terms of CFS framing process cost, CFS framing A1–A5 embodied carbon, and CFS framing process-hours using a fixed New Zealand case-study boundary and a model-calculated S2 reference configuration.

1.4.2: Objectives

1. Define a fixed CFS residential framing case-study boundary and model-calculated S2 reference configuration for comparing alternative framing configurations.
2. Formulate a transparent response model linking panel rationalisation intensity R_p, truss rationalisation intensity R_t, opening/detailing complexity OC, and prefabrication level P_f to CFS framing process cost, CFS framing A1–A5 embodied carbon, and CFS framing process-hours.
3. Implement an LHS-initialised NSGA-II optimisation workflow to generate nondominated CFS framing configuration alternatives within the defined case-study boundary.
4. Assess optimisation outputs using convergence diagnostics, S2-relative classification, sensitivity analysis, multi-seed stability checks, and Monte Carlo robustness testing under selected coefficient ranges.
5. Translate the optimisation, sensitivity, robustness, and ranking outputs into a dashboard-supported decision workflow for early-stage CFS framing interpretation.

1.4.3: Research Questions

The research questions follow from the integration gap identified in Section 1.3. If early-stage CFS framing configuration decisions affect process cost, A1–A5 carbon, and process-hours simultaneously, then the thesis must first define the relevant configuration variables, quantify their objective effects, identify attainable S2-relative Pareto alternatives, test whether selected alternatives remain defensible under uncertainty, and translate the results into a usable decision-support workflow. The research questions therefore move from formulation, to optimisation, to robustness, to practical interpretation.

- RQ1: How do panel rationalisation intensity, truss rationalisation intensity, opening/detailing complexity, and prefabrication level affect CFS framing process cost, CFS framing A1–A5 embodied carbon, and CFS framing process-hours within the fixed residential case-study boundary?
- RQ2: What Pareto-optimal cost–carbon–time configurations are achievable for the CFS residential framing case study relative to the S2 project-derived, model-calculated reference configuration?
- RQ3: How robust and practically defensible are selected Pareto configurations relative to S2 under uncertainty in cost rates, carbon factors, productivity assumptions, logistics, and waste allowances?
- RQ4: How can the optimisation and robustness results be translated into a practical dashboard-based decision-support workflow for early-stage CFS framing configuration decisions?

1.4.4: Scope and Boundary

This study is limited to a fixed single-storey residential CFS case-study building with a footprint area of 98.292 m² and an envelope/framing area of 217.2555 m². The model evaluates CFS framing configuration decisions within a controlled response-model and optimisation boundary. The included objectives are CFS framing process cost, CFS framing A1–A5 embodied carbon, and CFS framing process-hours.

The thesis also distinguishes between Pareto membership and improvement relative to S2. A configuration can be nondominated without improving all objectives relative to the reference case. Conversely, a subset of nondominated configurations may improve process cost, A1–A5 carbon, and process-hours simultaneously if the reference configuration is not locally efficient under the implemented response model. This does not remove the existence of trade-offs; rather, it makes the trade-off structure explicit by separating all-objective improvement alternatives from priority-specific alternatives. Chapter 5 therefore reports both the total nondominated archive and the subset of solutions that improve all three objectives relative to S2.

The study does not include whole-building architectural optimisation, whole-building life-cycle assessment, live AS/NZS 4600 member-capacity redesign, contractor tender pricing, or full construction scheduling. These exclusions define the intended modelling boundary rather than post-hoc limitations.

Table 1-1 Scope and boundary summary for the CFS framing optimisation framework.

Scope item	Exact boundary
Building geometry	Single-storey residential CFS case study; footprint fixed at 98.292 m2; envelope area fixed at 217.2555 m2.
Cost	CFS framing process cost only: steel material, factory processing, site labour, setup, transport, delivery/logistics, mobilisation, waste, and overhead.
Carbon	CFS framing A1–A5 only: A1–A3 steel, A4 transport, factory energy if included, A5 installation, and waste allowance.
Time	Construction process-hours only: factory production, setup, delivery/logistics, site installation, mobilisation, and weather allowance; not calendar duration.

Structural design	Optimisation uses pre-screened feasible sections and constraints; no live AS/NZS 4600 member-capacity iteration loop.
Optimisation method	NSGA-II with LHS initialisation
Sensitivity and uncertainty	OAT sensitivity and LHS-based screening using Spearman/PRCC are used; no variance-decomposition claim is made. Monte Carlo uncertainty analysis is applied to representative Pareto solutions using $N_{MC} = 1000$simulations under selected coefficient ranges.
Dashboard	Decision-support interface only; it reads exported CSV/PNG outputs and enables filtering, ranking, and comparison. It is not certified structural software.
Design-variable encoding	The optimiser controls the continuous variables R_p, R_t, OC, and P_f. Panel and truss scheme labels are derived reporting categories used for interpretation, plots, tables, and dashboard filtering.

The Monte Carlo analysis is interpreted as bound coefficient-uncertainty propagation, not as full model-form sensitivity or external empirical validation. These boundaries are used consistently when interpreting the optimisation results, S2 comparisons, uncertainty analysis, and dashboard outputs.

1.4.5: Limitations

This study has five main limitations. First, it is based on one fixed residential CFS case-study boundary, so the findings are case-study-specific and should not be generalised as universal CFS framing rules. Second, the optimisation does not perform live member-capacity checks, connection design, bracing verification, or iterative AS/NZS 4600 redesigns; it compares configuration alternatives after a fixed feasible structural framing system has been adopted. Third, the objectives are limited to CFS framing process cost, CFS framing A1–A5 carbon, and CFS framing process-hours, rather than whole-house cost, whole-building LCA, or full project calendar duration. Fourth, several coefficients are literature-supported, benchmark-based, or calibrated response-model assumptions rather than measured commercial project records. Fifth, the dashboard is a decision-support and communication layer for model outputs, not certified structural design, contractor pricing, carbon certification, or construction-programme software.

1.5: Study Contributions

Within the defined case-study boundary, this thesis makes five contributions to early-stage decision support for panelised CFS residential framing.

C1. CFS framing configuration response model.

The study formulates a transparent response model linking panel rationalisation intensity, truss rationalisation intensity, opening/detailing complexity, and prefabrication level to CFS framing process cost, CFS framing A1–A5 embodied carbon, and CFS framing process-hours.

C2. Project-derived S2-relative optimisation framework.

The study applies the response model to a real-project-derived CFS framing package and uses the model-calculated S2 configuration as a consistent reference for normalisation, improvement classification, and decision interpretation.

C3. LHS-initialised NSGA-II implementation with archive diagnostics.

The study implements an LHS-initialised NSGA-II workflow and evaluates the resulting nondominated archive using convergence, stability, feasibility, and S2-relative comparison metrics.

C4. Uncertainty-aware interpretation of representative solutions.

The study combines sensitivity analysis and Monte Carlo robustness testing to assess whether selected Pareto alternatives remain defensible under selected coefficient ranges.

C5. Dashboard-supported decision workflow.

The study translates optimisation, sensitivity, robustness, component-breakdown, and ranking outputs into a dashboard-supported workflow for early-stage CFS framing decision support.

Collectively, these contributions position the study as a bounded, coefficient-audited, uncertainty-aware decision-support framework for comparing CFS residential framing configurations. The contribution is not the invention of NSGA-II, Monte Carlo simulation, TOPSIS, or dashboard visualisation individually, but their traceable integration around a project-derived CFS framing case study and a clearly defined cost–carbon–process-hour decision problem.

1.6: Chapter Structure Summary

This study is organised into eight chapters and two supporting appendices. The structure follows the research workflow from problem definition and literature positioning, through model formulation and software implementation, to optimisation results, robustness testing, dashboard interpretation, and final conclusions. Table 1-2 summarises the role of each chapter and its relationship to the research objectives.

Table 1-2 Chapter structure mapped to research objectives.

Research component	Main purpose	Linked objective(s)	Main output
Chapter 1	Introduce problem, scope, aim, objectives, and contributions	O1-O5	Problem statement, scope boundary, aim, research questions
Chapter 2	Review CFS, prefabrication, carbon, optimisation, uncertainty, and dashboards	O1-O5	Literature review and research gap
Chapter 3	S2 baseline, $R_p/R_t/OC/P_f$ design-variable table, reporting-category definitions, objective functions, coefficient audit	O1-O5	S2 baseline, design-variable table, objective functions, coefficient audit
Chapter 4	Explain model implementation, verification, S2 reasonableness checks, and dashboard export traceability	O1-O5	LHS-NSGA-II workflow, continuous-variable encoding, data pipeline, dashboard architecture
Chapter 5	Present optimisation and Pareto-front approximation	O3-O4	Pareto front, representative solutions, S2 comparison, convergence results
Chapter 6	Present sensitivity, uncertainty, and robustness analysis	O4	OAT/LHS sensitivity, Spearman/PRCC screening for

			continuous variables, grouped reporting-category checks, Monte Carlo robustness
Chapter 7	Interpret results for practical decision support	O5	Practitioner decision example, dashboard interpretation, ranking discussion
Chapter 8	Conclude research and recommend future work	O1-O5	Answers to RQs, contributions, limitations, recommendations
Appendix A	Support reproducibility	O1-O5	Code/output mapping
Appendix B	Support dashboard use	O5	Dashboard user guide

Chapter 2: Literature Review and Research Gap

2.1: Literature Review Strategy

Chapter 2 reviews the literature required to support the CFS framing decision-support framework developed in this thesis. The review is organised around the engineering logic of the model rather than as a set of independent topic summaries. The central question is how project-derived panelised CFS residential framing alternatives can be screened in terms of process cost, A1–A5 embodied carbon, and process-hours when the relevant decisions involve rationalisation, detailing complexity, and prefabrication level.

The reviewed literature is grouped into four linked areas. The first concerns CFS framing systems and the fixed structural boundary, which defines what is being compared and what is not redesigned. The second concerns prefabrication, DfMA, cost, carbon, and productivity literature, which supports the modelling of configuration-dependent objective responses. The third concern's multi-objective optimisation, LHS initialisation, TOPSIS, and Monte Carlo uncertainty analysis, which support the search, ranking, and robustness-testing workflow. The fourth concerns visual analytics and dashboard-based decision support, which supports the translation of optimisation outputs into practitioner-facing interpretation.

NSGA-II, LHS, Monte Carlo simulation, TOPSIS, and dashboard visualisation are established methods. The literature gap addressed by this thesis is their transparent integration around a project-derived, New Zealand-context, CFS residential framing configuration-screening problem using panel rationalisation intensity, truss rationalisation intensity, opening/detailing complexity, and prefabrication level as the decision variables.

The chapter therefore moves from CFS framing and prefabrication mechanisms to objective modelling, optimisation and uncertainty methods, visual decision support, and finally the structured gap analysis.

The search strategy was structured to capture both engineering-domain literature and decision-method literature. This was necessary because the thesis is not only a CFS framing study and not only an optimisation study. It requires evidence for the CFS framing boundary, prefabrication and DfMA mechanisms, cost, carbon, process-hour response modelling, multi-objective optimisation, uncertainty testing, ranking, and dashboard-based communication.

Table 2-1 Literature search strategy.

Search area	Keywords used	Purpose
CFS / LGSF framing	cold-formed steel, light gauge steel, LGSF, light steel framing	Identify CFS residential and structural framing studies
CFS optimisation	cold-formed steel AND optimisation, genetic algorithm, multi-objective	Identify whether CFS optimisation is member-level or system-level
Prefabrication	prefabrication, off-site construction, panelised, modular, DfMA	Review production, standardisation, waste, cost, and time studies
Cost–carbon–time optimisation	cost, time, carbon, NSGA-II, multi-objective, construction	Identify studies with simultaneous objective functions
Embodied carbon	embodied carbon, upfront carbon, A1–A5, LCA, steel	Support the carbon boundary and emissions assumptions

Search area	Keywords used	Purpose
Uncertainty and sensitivity	Latin Hypercube Sampling, Monte Carlo, PRCC, Spearman	Support robustness and sensitivity methods
Dashboard / decision support	dashboard, decision support, visual analytics, Pareto visualisation	Support the dashboard and interpretation layer

The search was therefore used to test whether a directly comparable integrated framework already existed. Individual methods and domains were well represented in the literature, but the review did not identify a study that combined project-derived CFS residential framing data, CFS-specific configuration variables, cost–carbon–process-hour optimisation, uncertainty testing, stakeholder ranking, and dashboard-supported interpretation within one transparent New Zealand case-study workflow.

The literature review therefore supports the research at two levels: first, by identifying sources that justify each individual modelling component, and second, by showing that these components are not commonly integrated into one CFS residential framing decision-support framework.

The search was conducted using thematic search areas rather than one combined query because the research integrates CFS framing, prefabrication, optimisation, uncertainty, and dashboard decision support. A single combined search would risk excluding relevant studies that address only one part of the framework. The selected studies are therefore synthesised through a feature-based gap matrix rather than presented as a paper-by-paper catalogue.

Each section of this literature review is structured to justify a specific component of the research framework. Sections 2.2 and 2.3 establish the CFS structural boundary and prefabrication context. Sections 2.4 to 2.6 justify the three objective-function domains: CFS framing process cost, CFS framing A1–A5 carbon, and CFS framing process-hours. Sections 2.7 to 2.9 justify the optimisation, sampling, sensitivity, and uncertainty methods, while Section 2.10 positions the dashboard as a visual decision-support layer. Section 2.11 then synthesises the reviewed literature into the integration gap addressed by this study.

The purpose of this structure is to connect each literature domain to a modelling decision in the thesis. The detailed integration gap is synthesised at the end of the chapter in Section 2.11.

2.2: Cold-Formed Steel Structural Systems

Cold-formed steel (CFS) structural systems are formed by shaping thin steel sheets into structural profiles such as lipped channels, tracks, studs, joists, and truss members at ambient temperature. Their structural efficiency comes from using thin-walled profiles to provide high strength-to-weight performance, dimensional accuracy, and compatibility with repetitive framing systems. In Australia and New Zealand, the structural design context is governed by AS/NZS 4600:2018, which sets out minimum requirements for structural members cold-formed from carbon or low-alloy steel sheet, strip, plate, or bar not more than 25 mm thick and used for load-carrying purposes in buildings (Standards Australia / Standards New Zealand, 2018). The North American AISI S100 specification provides a comparable international design basis for load-carrying CFS members and reflects continuing research into CFS member behaviour. For this study, CFS is therefore treated as a structurally established framing system, but the optimisation does not attempt to redesign the structural member capacities (American Iron and Steel Institute, 2022; Hancock et al., 2026).

A key characteristic of CFS members is that thin-walled behaviour makes stability effects central to structural design. Local, distortional, and global buckling may govern member resistance, and established CFS design procedures such as effective-width and direct-strength methods are used to address these modes. This structural behaviour is important background for the thesis because it explains why member sizing, section geometry, connection design, bracing design, and code-level structural verification are treated as fixed boundary conditions rather than optimisation variables (Schafer, 2008; Yu et al., 2020). (Anil Kumar & Kalyanaraman, 2018). (Standards Australia / Standards New Zealand, 2018; Yu et al., 2020).

For the case-study model, the structural section is fixed as F325iT Imported Section, 89 S 41 / 0.75 / G500 / Z275, with a steel thickness of 0.75 mm. Because this section is fixed, the optimisation does not vary steel thickness or member capacity. The literature reviewed in this section therefore supports the structural boundary of the study, while the optimisation itself focuses on CFS framing configuration, production, logistics, embodied carbon, and process-hour effects. The current model treats steel thickness, grade, and coating class as fixed project boundary conditions, and the optimisation is restricted to four code-aligned configuration variables: panel rationalisation intensity R_p, truss rationalisation intensity R_t, opening/detailing complexity OC, and prefabrication level P_f. Panel and truss scheme labels are derived from $R_p$and $R_t$only for reporting and interpretation. This fixed-section assumption allows the study to isolate production-configuration effects while leaving detailed member-capacity verification outside the literature-review argument. The objective is to evaluate how production-oriented decisions influence CFS framing process cost, CFS framing A1–A5 carbon, and construction process-hours within a fixed structural boundary.

Table 2-2 CFS section and structural boundary summary.

Item	Value / description	Unit	Source/ basis	Role in study
CFS section source	F325iT Imported Section	—	Project material schedule / Frame CAD section listing	Fixed case-study section
Section designation	89 S 41	—	Project material schedule	Fixed structural boundary
Steel thickness	0.75	mm	Project material schedule	Fixed; not optimised
Steel grade	G500	—	Project material schedule	Fixed material grade
Coating class	Z275	—	Project material schedule	Fixed corrosion-protection assumption
Wall panels	48	count	Case-study panel schedule	Used in panel scheme and repetition logic
Roof trusses	58	count	Case-study truss schedule	Used in truss scheme and repetition logic
Total steel mass	2195.264	kg	Case-study quantity take-off / model baseline	Used in cost and carbon objectives

Item	Value / description	Unit	Source/ basis	Role in study
Structural capacity treatment	Not re-solved inside optimiser	—	Research scope boundary	Optimisation focuses on configuration, not member capacity
Optimised variables	R_P, R_t, OC, and P_f	Dimensionless / fraction	Implemented optimisation model	Continuous CFS framing rationalisation and prefabrication variables
Reporting categories	panel_scheme, truss_scheme	Index	Derived from $R_p$and R_t	Used for plots, tables, dashboard filters, and practitioner interpretation

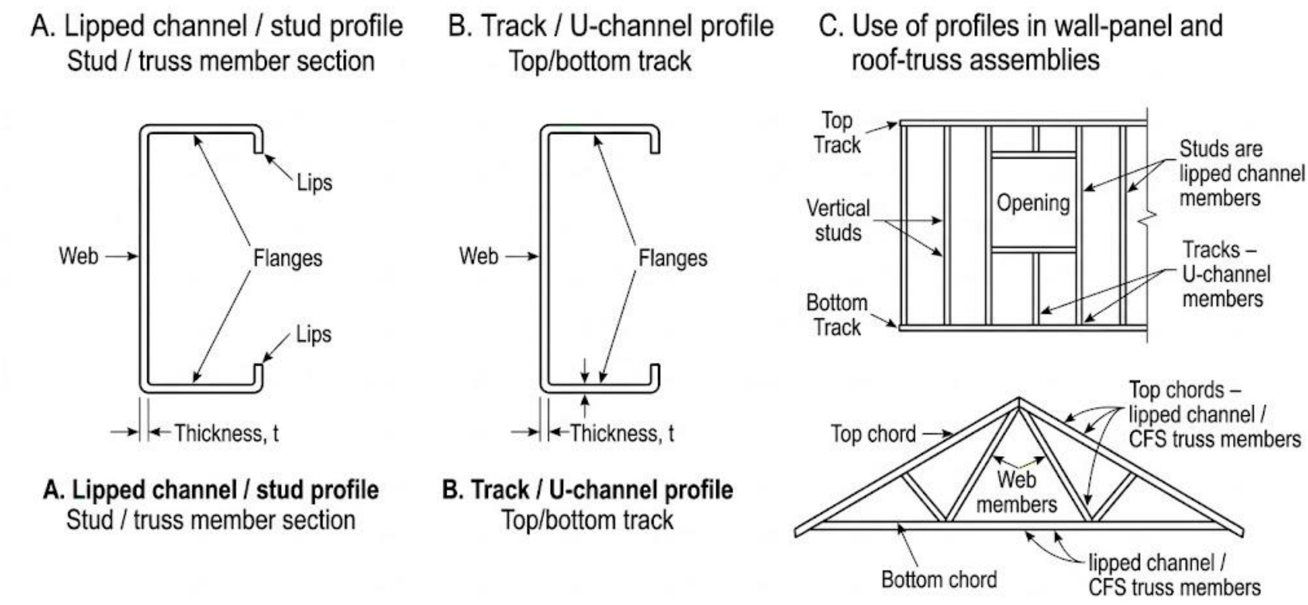

Figure 2-1 Typical CFS framing profiles and their use in wall-panel and roof-truss assemblies. (Source: Author's own schematic, adapted from Australian Steel Institute CFS section-profile guidance and SCI P402.)

Previous CFS optimisation research has mainly focused on member-level or section-level optimisation, including cross-sectional shape, thickness, local buckling behaviour, and thermal or structural performance. (Liang et al., 2022) review CFS optimisation studies and show that the dominant focus remains on structural performance of CFS sections and systems rather than production-level configuration decisions. This supports the positioning of the present research: the fixed section boundary is not a modelling weakness, but a deliberate scope decision that allows the study to isolate panelisation, repetition, opening/detailing complexity, and prefabrication level as production-oriented variables.

The New Zealand CFS framing context is also shaped by the availability of digital design-to-manufacture workflows. FRAMECAD-based systems are used in New Zealand light-gauge-steel framing practice to connect framing design, detailing, roll-forming, panel fabrication, and truss manufacture. Manufacturer and fabricator information shows that CFS framing is commonly delivered through a digitally coordinated production chain rather than only as loose structural members. This is relevant to the present study because the optimisation variables are not member-capacity variables; they represent production-configuration decisions, including panel rationalisation, truss rationalisation, opening/detailing complexity, and prefabrication level. The manufacturer literature is used here only to define the practical production context and does not replace project-specific validation, structural design verification, or supplier quotation data (*FRAMECAD*, 2026; *Steel Frames NZ*, 2026).

The CFS literature therefore supports treating the structural system and framing section as a fixed boundary for this thesis. This is important because the optimisation problem is not member sizing, connection design, bracing design, or code-based structural redesign. Instead, the fixed CFS framing package provides the physical basis for comparing production-configuration decisions: panel rationalisation, truss rationalisation, opening/detailing complexity, and prefabrication level. The detailed S2 case-study boundary and design-variable encoding are defined in Chapter 3.

This structural boundary is necessary because the optimisation compares CFS framing configuration decisions, not alternative structural systems or live member-capacity designs.

2.3: CFS Prefabrication and Panelised Construction

Prefabrication in building construction may be viewed as the movement of construction activities from the project site to a controlled production site. It can be found in a range of prefabricated components, panelised systems or complete volumetric/modular units. For the residential light steel construction, the most relevant category for this study is the panelised prefabrication, which is the manufacture of wall panels, roof trusses, floor cassettes or other components prior to delivery to site. SCI advises that light steel systems involve the use of galvanised cold-formed steel sections as the main structural components and that these sections may be prefabricated into panels. Extensive prefabrication literature also reveals that off-site production needs to be considered as a whole and not just based on the direct cost, as it can be beneficial based on labour, logistics, quality, programme and production effects. As a result, prefabrication is considered not merely as a construction technique, but as a quantifiable design and production choice, in this study (Li et al., 2014; Mao et al., 2016).

The typical sequence of a CFS panelised system is from the start of digital modelling, member scheduling, roll-forming, cutting and punching, labelling, through to the assembly of the panels within the factory, the transport to the site, handling and installation, as well as the connection to the adjoining panels and structural elements. This workflow differs from purely site-built construction, as some labour and quality-control aspects are transferred to the factory, while there is an increased need for co-ordination, transport and installation planning. In their study, Pan et al. looked at the adoption of off-site techniques in leading housebuilders in the UK and found that the benefits and constraints of implementing off-site play a role in their decision (Pan et al., 2008). Likewise, the New Zealand off-site construction environment indicates that some benefits of modular and off-site construction (schedule and productivity) are limited by the industry's ability, procurement, supply chain coordination, and barriers to implementation on a project (Shahzad et al., 2024). This is relevant to the present research because the model does not assume that increasing prefabrication is always beneficial. Instead, it considers the impact of prefabrication on the factory, site, logistics, costs, carbon and process time.

Standardisation and repetition are important principles in prefabricated construction because they can reduce production variability, simplify assembly processes, and improve process coordination. In the New Zealand context, BRANZ SR312 is used in this review only as contextual support for the broader benefits and constraints of prefabrication, including quicker construction, improved quality, reduced site disruption, and standardisation opportunities. It is not used as direct evidence for a numerical CFS learning-curve coefficient. The learning-curve mechanism is instead supported more directly by repetitive-construction and modular-manufacturing literature, which shows that repeated tasks can improve productivity but that learning effects may fluctuate rather than follow a constant linear pattern. Therefore, in this study, repetition and standardisation of CFS panels and trusses are treated as bounded process-efficiency assumptions within the response model, not as universal productivity laws (Ian Page & David Norman, 2024).

The learning-curve mechanism is more directly supported by Lu et al. (2023), who showed that accumulated experience in modular construction manufacturing can improve productivity, although the learning rate may fluctuate rather than follow a strictly linear or constant pattern (Lu et al., 2023). Similarly, Widanage and Kim (2024) support the role of DfMA–BIM integration in reducing assembly time and site labour, but this evidence treated as modular/DfMA evidence rather than direct proof for CFS framing (Widanage & Kim, 2024). Therefore, in this study, repetition and standardisation of CFS panels and trusses are treated as bounded process-efficiency assumptions, supported by prefabrication, modular manufacturing, and DfMA literature, rather than as guaranteed universal productivity rules. In this study, panel and truss rationalisation are represented by the continuous variables $R_p$and R_t. Their effects on cost and time are represented through setup, factory processing, and installation productivity terms, while their carbon effects are indirect through steel mass, waste, and logistics assumptions.

Panelised CFS framing is a process of standardisation and repetition at the production level. It can shorten setup time, simplify CNC programming, accelerate the learning curve in production, and facilitate site installation, which can be repeated in the panel width and truss span families. However, excessive rationalisation may introduce oversizing, trimming, additional coordination, or architectural constraints (Lu et al., 2021). This tension is in line with the literature on design-for-manufacture-and-assembly (DFMA), which focuses on modularity, standardization, simplified assembly, and early design coordination, and that manufacturing and assembly constraints taken into account in the design process (Li et al., 2014). In this study, this production logic is represented using panel rationalisation intensity $R_p$and truss rationalisation intensity R_t. These continuous variables represent progressive standardisation of panel-width and truss-span families relative to the S2 reference condition. For communication and dashboard filtering, the resulting $R_p$and R_t values are later grouped into derived reporting labels: panel_scheme and truss_scheme.

Opening and detailing complexity is treated as a design-for-manufacture variable rather than an architectural opening variable. In panelised CFS construction, openings require lintels or headers, jamb studs, trimmers, local reinforcement, service penetrations, cutting operations, and site fit-up coordination. These details affect production complexity even when the number of windows and doors remains unchanged. Therefore, this study does not use window-to-envelope ratio as an optimiser variable and does not add or remove openings. Instead, opening/detailing complexity (OC) represents the degree of standardisation in lintels, jambs, trimmers, service penetrations, CNC cutting, and site fit-up. This is aligned with DfMA thinking because the variable captures how easily a panel can be manufactured and assembled, rather than changing the architectural design intent. In the model, the reference number of openings is fixed at $N_{openings}$ =12, and OC ranges from 0.80 to 1.30, where lower values represent more standardised detailing and higher values represent more customised detailing. Window-to-envelope ratio is therefore not treated as an optimiser variable in this study; the number of openings remains fixed, and only the complexity of their CFS detailing is represented.

The prefabrication and DfMA literature show that standardisation, repetition, factory/site coordination, and assembly simplification can influence construction performance. However, most studies treat these issues as general off-site-construction benefits, adoption barriers, or qualitative DfMA principles rather than as linked CFS framing response variables. They rarely quantify panel rationalisation, truss rationalisation, opening/detailing complexity, and prefabrication level together in one cost–carbon–process-hour optimisation framework. This matters because CFS framing decisions are not controlled by prefabrication level alone; the benefits of off-site production depend on how repeatable, standardised, transportable, and buildable the framing configuration is.

The prefabrication level (P_f) represents the relative shift of CFS framing work from site-based construction to factory-based production. In the model, P_f ranges from 0.50 to 0.90, with three scenario anchors: S1 low prefabrication $(P_f = 0.50)$, S2 medium prefabrication $(P_f = 0.72)$, and S3 high prefabrication $(P_f = 0.90)$. Increasing prefabrication can reduce site labour, weather exposure, and construction-stage waste, but it may also increase factory processing, logistics, transport, handling, and design coordination. Jaillon et al. show that prefabrication can reduce construction waste in building projects, while Mao et al (Jaillon et al., 2009; Mao et al., 2016). show that off-site construction can introduce additional cost items through production, transportation, and design coordination. Ferdous et al. also emphasise that modular construction can improve speed, quality, predictability, and waste performance, but adoption is still affected by technical and industry constraints (Ferdous et al., 2019). For this reason, P_f is modelled as a trade-off variable rather than as a fixed assumption that higher prefabrication is always better.

Table 2-3 Prefabrication metrics and relevance across related studies.

Study	Context	Main metric / focus	Cost	Carbon / waste	Time /	Standardisatio	Relevance to this study
SCI P402 / Yandzio et al.	Light steel residential framing	Light steel panels and residential framing workflow	○	–	○	○	Supports CFS panelised framing context.
Blismas et al. (2006)	Off-site production	Holistic benefit evaluation	✓	○	✓	○	Supports evaluating prefabrication beyond direct cost.
Pan et al. (2008)	Housebuilders and off-site construction	Adoption drivers/barriers	○	–	○	○	Supports practical implementation and adoption constraints.
Jaillon et al. (2009)	Building prefabrication	Waste reduction potential	○	✓	○	○	Supports waste and construction-stage impact discussion.
Li et al. (2014)	Prefabricated construction management	Management research review	○	○	○	○	Supports coordination and production-management framing.
Mao et al. (2016)	Sustainable off-site construction	Cost structure and case studies	✓	○	–	○	Supports additional production, transport, and design-cost discussion.
Lu et al. (2021)	DfMA in construction	Manufacturability and assembly	○	○	○	✓	Supports panel/truss rationalisation and opening/detailing standardisation.
Ferdous et al. (2019)	Modular buildings review	Benefits, constraints, and opportunities	○	○	✓	○	Supports speed, predictability, quality, and adoption constraints.
Shahzad et al. (2024)	New Zealand modular off-site construction	Benefits, constraints, enablers	○	○	✓	○	Provides NZ-specific off-site construction context.
Wang et al. (2024)	Prefabricated component optimisation	Cost-duration-carbon optimisation	✓	✓	✓	○	Closest optimisation comparison, but not CFS panel/truss/opening specific.
this study	Panelised CFS residential framing	R_p, R_t, OC, P_f with panel_scheme and truss_scheme used as derived reporting categories.	✓	✓	✓	✓	Integrates CFS configuration variables with multi-objective optimisation and decision support.

Legend: *✓ = explicitly addressed; ○ = indirectly or partially addressed; – = not a main focus.*

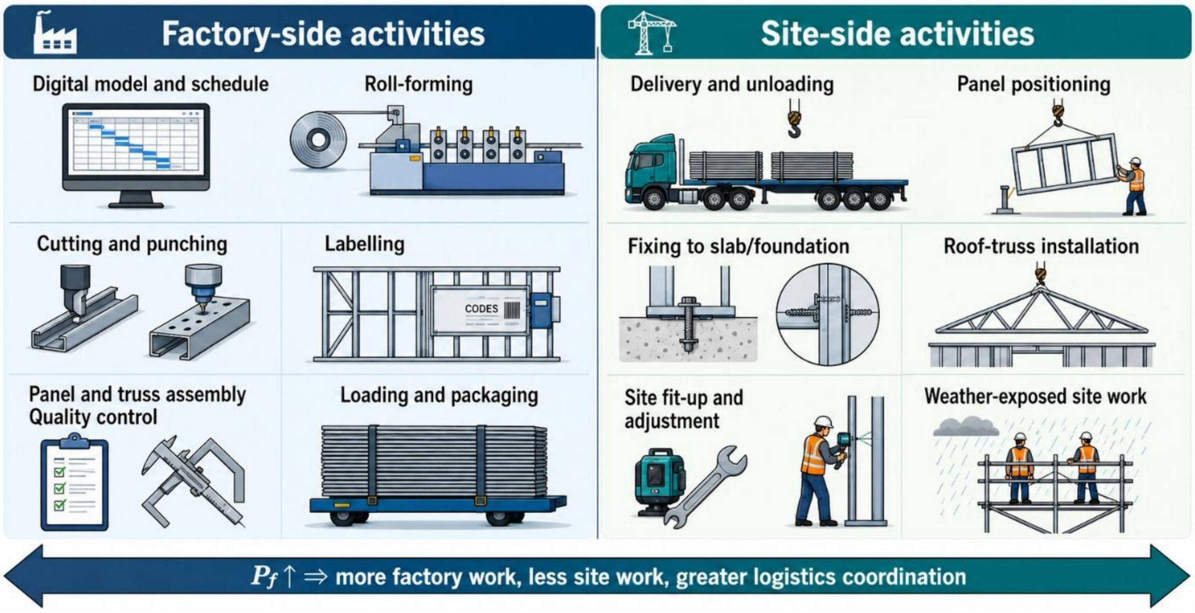

Figure 2-2 Conceptual split between factory-side and site-side activities in panelised CFS prefabrication. (Source: Author's own schematic, informed by light steel framing, off-site production, and DfMA literature.)

The prefabrication and DfMA literature is directly relevant to the formulation of R_p, R_t, OC, and P_f. Panel and truss rationalisation represent the production-side benefit of repetition and reduced variation. Opening/detailing complexity represents the fabrication and site-fit-up burden associated with non-standard openings, lintels, jambs, trimmers, and local details. Prefabrication level represents the shift of work from site installation to factory preparation. These variables are therefore not arbitrary optimisation inputs; they are response-model proxies for documented production and constructability mechanisms.

2.4: Construction Cost and CFS Framing Cost Modelling

Cost estimation for CFS framing can be undertaken at several levels, ranging from whole-building cost comparison to component-level framing assessment. Because this thesis evaluates early-stage CFS framing configuration decisions, the relevant cost boundary is CFS framing process cost rather than whole-house construction cost. The reviewed literature supports the use of cost as a comparative design metric, but it also shows that the cost boundary must be clearly stated because results change depending on whether the analysis includes framing only, the full building, or life-cycle costs (Doctolero & Batikha, 2018; N. Usefi et al., 2021)

For prefabricated CFS framing, cost is not governed by steel mass alone. Factory processing, setup, labour, transport, delivery logistics, site mobilisation, opening/detailing effort, waste handling, and overhead can all influence the process-cost response. Off-site construction cost studies are therefore useful because they show that prefabrication shifts cost between factory, logistics, and site activities rather than simply reducing cost. This supports the component-based cost structure used later in the thesis, where the CFS framing process-cost objective is decomposed into traceable cost terms instead of represented by a single global cost-per-square-metre rate (AbouHamad & Abu-Hamd, 2019; Mao et al., 2016)

Recent studies on prefabricated buildings optimisation also confirm the importance of considering cost, time, carbon and uncertainty. Under uncertainty, Yuan et al. proposed a simulation and optimisation framework for prefabricated building construction, which explored how the interacting activity logic and multiple objectives influence the performance of prefabricated construction (Yuan et al., 2024). Song et al. presented a multi-objective optimisation model of assembled buildings by incorporating carbon emissions, duration as well as cost and quality level as project objectives (H. Song et al., 2024). Xu et al. also proposed a multi-objective optimisation model for prefabricated buildings with the objective functions of cost, duration and carbon emissions, and the construction technology and prefabrication rate were also included in the decision structure (Xu et al., 2025). These studies are not CFS panel/truss configuration studies, but they are significant in that they reveal that cost estimation in prefabricated construction is related to the construction process and assessed in conjunction with carbon and time rather than as an independent measure.

Another problem is that when comparing costs in a whole building, the contribution of the structural framing system is obscure. When comparing the cost of steel structures to other construction materials it is important to keep in mind that numerous factors outside of the structural frame contribute to the overall cost of the project, such as services, fit-out, and allowances, which make cost comparisons more difficult to see. Industry cost guides are used only as secondary context, because whole-building cost comparisons include many non-framing items (Brown et al., 2020; *Cost Comparison Studies*, 2026; *NZ Construction Costs Update - QV CostBuilder*, 2025). This is directly relevant to the present research since optimisation is not aimed at estimating the total house cost. Rather, it separates the cost of the CFS framing process to allow panel rationalisation intensity, truss rationalisation intensity, opening/detailing complexity, and prefabrication level to be studied without adding unrelated elements like foundations, finishes, services, kitchens, baths, or operating energy systems. Any cost per square metre quoted thereafter in this study should therefore be regarded as only an indicator of the CFS framing scope and not a cost for the whole house construction.

In line with this, the cost model presented in this study is organised by dimensionally consistent costs: steel material costs (NZD/kg), factory processing costs (NZD/kg), site labour costs (NZD/hr), set up costs (NZD/type), transport costs (NZD/t-km), waste handling costs (NZD/t), and overhead (percentage allowance). This makes the model transparent and auditable as it is based on physical, production activities. In addition, it fits into the overall literature which suggests that material quantity, factory production, site labour, transport, logistics and uncertainty are important trade-offs in making prefabrication decisions. The cost model is thus not shown as a market quotation or quantity-surveyor estimate, but as a comparative process-cost model for cost comparisons of CFS framing alternatives within a fixed case-study perimeter. The cost model provides a consistent process-cost response for comparing alternatives under the S2-calibrated case-study boundary.

Table 2-4 Cost-model components and literature support.

Cost component in this study	Unit	Why it is included	Literature / evidence support
Steel material	NZD/kg	Captures framing mass effect and material-cost sensitivity	CFS comparative studies commonly report structural weight and material cost as key comparison metrics.
Factory processing	NZD/kg	Represents roll-forming, cutting, punching, handling, and assembly effort	Off-site construction studies show that prefabrication shifts work into production/manufacturing stages.
Site labour	NZD/hr	Captures installation effort affected by prefabrication level and detailing complexity	Prefabrication literature links off-site production to changes in site labour and productivity.
Setup / CNC / jigging	NZD/type	Captures cost of unique panel and truss families	DfMA and prefabrication literature emphasise standardisation, repetition, and setup reduction.
Transport	NZD/t-km	Captures effect of moving prefabricated CFS components to site	Off-site construction cost studies identify transportation as a relevant cost category.
Delivery / logistics	NZD/hr or allowance	Captures coordination, unloading, sequencing, and site-handling burden	Prefabricated construction studies show logistics and coordination can offset some site-efficiency gains.
Waste handling	NZD/t	Captures offcuts and disposal/handling implications	Prefabrication studies evaluate waste reduction and material-efficiency impacts.
Overhead	% of direct cost	Represents coordination, management, and project-level indirect allowance within framing scope	Construction cost frameworks commonly separate direct and indirect/project overhead costs.

The cost literature therefore supports a component-level CFS framing process-cost objective, but it does not provide a directly transferable New Zealand CFS residential framing cost model for the exact R_p, R_t, OC, and P_f decision variables used in this study.

The cost coefficients used later in the methodology are not introduced as universal market rates or tender prices. Instead, they are treated as case-study calibration coefficients for comparing CFS framing alternatives under a consistent process-cost boundary. The literature supports the inclusion of material, factory processing, site labour, setup, transport, waste, and overhead components, while the numerical coefficient values are defined and audited in Chapter 3 using project-derived data, supplier or benchmark information, and uncertainty ranges. This separation is important because Chapter 2 justifies the cost-model structure, whereas Chapter 3 defines the actual model coefficients used in the optimisation. In the current model, setup and rationalisation effects are not treated as literally discrete panel or truss choices. Instead, they are represented through R_p and R_t, which influence effective panel-class and truss-type response quantities used in the cost model.

A limitation of the available literature is the scarcity of publicly available New Zealand CFS framing cost data at the level of panel classes, truss families, opening/detailing complexity, and prefabrication level. Manufacturer and fabricator information confirms that digital design, roll-forming, panel fabrication, and truss manufacture are established in New Zealand practice, but it does not usually disclose project-level cost rates, factory productivity rates, or setup-cost structures. For this reason, the present study does not claim that the cost model represents a market-price estimator. The literature review supports the cost-component structure, while Chapter 3 defines the numerical coefficients as case-study response-model assumptions with uncertainty ranges.

The cost literature supports a component-level process-cost objective rather than a whole-house tender-price objective. This distinction is important because the thesis compares CFS framing configuration alternatives, not complete building procurement options. The cost objective must therefore remain sensitive to steel mass, factory processing, setup, transport, waste, site labour, and prefabrication-related shifts in effort. A single building-level $/m² cost rate would hide the mechanisms that rationalisation, opening/detailing complexity, and prefabrication level are intended to test.

The cost literature therefore supports a component-level CFS framing process-cost objective because whole-house cost rates would obscure the mechanisms affected by rationalisation, detailing, logistics, factory processing, and site labour.

2.5: Embodied Carbon and A1–A5 Framing Boundary

Life-cycle assessment provides the methodological basis for estimating embodied carbon in building systems. Life cycle assessment (LCA) principles and framework that are defined in ISO 14040: Goal and scope definition, Inventory analysis, Impact assessment, Interpretation, Reporting, Limitations. The requirements and guidelines for conducting and reporting LCA studies are given in ISO 14044 (ISO 14040:2006, 2006; ISO 14044:2006, 2006). EN 15978 for buildings is a structure for building-level assessment comprising of product, construction, use, end-of-life and beyond-boundary modules. Additionally, MBIE's whole-of-life embodied carbon framework indicates that embodied carbon calculations for buildings are based on EN 15804, EN 15978 and ISO 21930 (European Committee for Standardization, 2011; Ministry of Business, Innovation and Employment, 2020b). In this study the carbon boundary is deliberately limited to CFS framing A1–A5 carbon, since the research question is about alternative prefabricated CFS framing configurations, rather than operational energy, maintenance, replacement, deconstruction, or end of life recovery.

The A1–A3 module represents the product stage, including raw-material supply, transport to the manufacturing site, and manufacture of the construction product. For steel framing systems, this stage is strongly controlled by the steel emission factor and the mass of steel included in the framing package. However, a single generic steel carbon factor is not technically equivalent to a product-specific value for thin-gauge galvanised CFS. Steel embodied-carbon values vary with production route, recycled content, galvanising or coating system, product thickness, country of manufacture, allocation method, and whether the source is a generic database value or a verified environmental product declaration. Therefore, ICE database values, World Steel data, and steel EPDs are treated in this study as benchmark sources rather than interchangeable values (Inventory of Carbon and Energ, 2026; World Steel Association, 2022).

The central A1–A3 steel emission factor used in this study is therefore treated as a benchmark coefficient rather than as a product-specific EPD value. This is appropriate for early-stage decision support, provided that the coefficient is clearly reported and sensitivity-tested. The ICE database provides general embodied-carbon factors for construction materials, while World Steel Association life-cycle inventory studies show that steel emissions vary substantially with production route, energy source, scrap content, and processing pathway (Inventory of Carbon and Energ, 2026; New Zealand Green Building Council, 2024; World Steel Association, 2022). Therefore, the selected CFS steel factor is used as a transparent modelling assumption and is tested within lower and upper uncertainty bounds rather than being claimed as a certified manufacturer-specific value.

The A4 module covers the transport of construction products to the construction site. A freight emission factor and transport distance, along with CFS framing mass transported, are used in this study to express A4 transport (*Life Cycle Stages | One Click LCA*, 2025). The default transport distance is 121 km, the default freight emission factor is 0.135 $kgCO_{2e}$/t-km, and the selected transport scenario has an equivalent value of 0.016335 $kgCO_{2e}$/kg (Vickers et al., 2024). It is important to note that this is a scenario assumption and not a measured field value and is applicable to the assumed New Zealand supply chain (Onat & Kucukvar, 2020). The A4 result is thus location sensitive: a project fabricated near the site would have a lower A4 transport contribution, whilst a project that needs a longer freight haul, more transport legs or heavier components would have a higher A4 transport contribution. Therefore, the model term A4 is given as a transparent model term, instead of being embedded in a single global carbon factor.

A5 emissions include emissions during construction (installation) and construction waste. This can include site handling, lifting, temporary power, crew activity, frame erection waste and more for CFS framing. The NZGBC embodied carbon methodology considers upfront carbon as embodied in the manufacture of products (A1–A3), transport to site (A4), construction (A5) and construction waste. The envelope area is calculated in the current model as the area of the building's envelope, and the A5 installation is represented by a CFS-framing-scope installation allowance of 3.20 $kgCO_{2e}/m^2$ (New Zealand Green Building Council, 2024). Factory energy is modelled separately where included. Site energy is not added separately when the A5 allowance already includes equipment, fuel, waste, or site activity effects, because this would risk double counting construction-stage impacts (Pomponi & Moncaster, 2018)

Waste effects have been incorporated as prefabrication can have an impact on offcuts, rework, site disturbance and material handling. The potential for construction waste reduction as a result of prefabrication has been established through studies conducted in the field, but this varies depending on the type of construction project, design principles for standardisation and quality of implementation (BRANZ, 2025; Jaillon et al., 2009). Waste carbon in this study is considered as a steel waste allowance, which is related to the steel mass and steel emission factor (Ministry of Business, Innovation and Employment, 2020). The current model assumes a reference waste steel rate of 5% and takes offcuts waste carbon into account if the waste-carbon switch is turned on. This is clearly differentiated from A5 construction waste disposal. Waste may be regarded as a demand on additional material and would be in the product burden group A1–A3, or as a handling or disposal activity on site, it would be in A5. The model must then clearly define the relationship between waste and the representation of waste, to prevent combining fabrication waste, site waste and disposal emissions without being able to account for them separately.

Therefore, this study's carbon objective is a scope-limited upfront carbon model for CFS framing, not a building LCA. Operational energy and water use, maintenance and repair, replacement and refurbishment, deconstruction, end-of-life processing and disposal, and module D recovery or reuse credits are excluded. The exclusions are in line with the research aim as they are comparing different CFS framing configurations within a consistent boundary from the cradle to site. Emission factors and scenario assumptions used in the model are not substitutes for product-specific EPDs, full building LCA, or third-party carbon certification and will be used to assist in early-stage decision-making. This limitation is acceptable provided the boundary is stated clearly and applied consistently across all alternatives. The model estimates comparative CFS framing upfront carbon under a consistent A1–A5 boundary.

$$E_{\text{CFS,A1-A5}} = E_{\text{A1-A3}} + E_{\text{A4}} + E_{\text{fac}} + E_{\text{A5}} + E_{\text{waste}}$$

Equation 2-1

$$E_{\text{A1-A3}} = W_{\text{steel}} \cdot EF_{\text{steel}}$$

Equation 2-2

$$E_{\text{A4}} = \left(\frac{W_{\text{steel}}}{1000}\right) d_{\text{transport}} EF_{\text{freight}}$$

Equation 2-3

$$E_{\text{A5}} = A_e \cdot EF_{\text{A5}}$$

Equation 2-4

The equation represents the CFS framing carbon boundary only. It is not a whole-building embodied-carbon equation. Factory energy and waste are retained as explicit model terms so that they can be included, excluded, or sensitivity-tested without hiding them inside a single global carbon factor.

Table 2-5 Carbon benchmark review for CFS framing A1–A5 boundary.

Carbon source	**Module**	**Typical benchmark / unit**	**Key source type**	**Relevance to this study**
Life-cycle boundary	A1–A5, B, C, D	EN 15978 building life-cycle modules	ISO 14040/14044; EN 15978; MBIE; NZGBC	Supports limiting the model to CFS framing A1–A5 carbon and excluding B, C, and D.
Steel production	A1–A3	$kgCO_{2e}$/kg steel	ICE database; World Steel	Main carbon driver because CFS framing carbon is strongly controlled by total steel mass and steel emission factor.

Carbon source	**Module**	**Typical benchmark / unit**	**Key source type**	**Relevance to this study**
			Association; steel EPDs	
Cold-formed / light-gauge steel EPDs	A1–A3	approx. $kgCO_{2e}/kg$ product	Product-specific or industry-average EPDs	Provides benchmark values for checking whether the selected steel carbon factor is reasonable.
Freight transport	A4	$kgCO_{2e}/t\text{-}km$	NZ freight-emission guidance / transport studies	Supports modelling transport as a scenario term based on steel mass, distance, and tonne-kilometre factor.
Factory energy	A3 / factory process	$kgCO_{2e}/kWh$	NZ electricity/grid emission factors	Relevant if factory electricity is modelled separately; must avoid double counting with product-stage factors.
Installation and site activity	A5	$kgCO_{2e}/m^2$ or activity-based allowance	NZGBC / BRANZ / LCA guidance	Supports inclusion of CFS framing installation, handling, equipment, and construction-stage activity.
Waste and offcuts	A1–A3 or A5	waste rate or $kgCO_{2e}$ allowance	NZGBC / BRANZ / prefabrication literature	Waste must be clearly assigned either as material uplift or construction-stage waste to avoid double counting.
End-of-life / recovery credit	C / D	excluded	EN 15978 boundary logic	Excluded because this study compares upfront CFS framing configurations, not whole-life recovery scenarios.

Table 2-5 summarizes the main benchmark sources relevant to the CFS framing A1–A5 carbon boundary used in this study. The table distinguishes between product-stage steel emissions, freight transport, factory or site energy, installation, and waste effects. This is necessary because the carbon objective is a comparative CFS framing carbon model, not a whole-building life-cycle assessment.

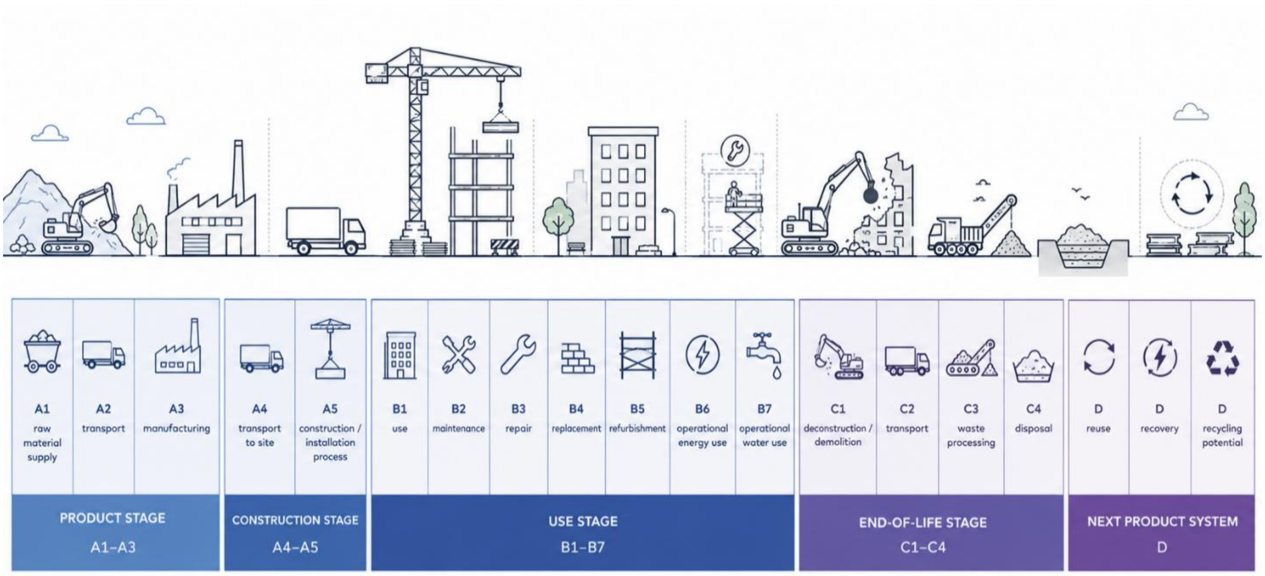

Figure 2-3 Stages of the building life cycle according to EN 15978.

The carbon model used in this study estimates comparative CFS framing upfront carbon under a consistent A1–A5 boundary. The reviewed carbon literature therefore supports an upfront A1–A5 framing-level boundary for comparing alternatives under a consistent scope. This boundary is appropriate because the research compares alternative CFS framing configurations within the same case-study geometry, rather than assessing the full environmental performance of the dwelling.

The need for framing-level carbon information is also linked to New Zealand's emerging building-carbon policy and procurement context. MBIE's Building for Climate Change programme and whole-of-life embodied-carbon framework identify the reduction and reporting of building-related embodied carbon as part of the broader transition toward lower-emission buildings. Procurement guidance also indicates that carbon-emission reduction can influence supplier assessment and tender-selection processes. This policy and procurement context supports the practical relevance of early-stage CFS framing carbon comparison, even though the present study remains limited to a case-study CFS framing A1–A5 boundary rather than a certified whole-building LCA (*Procurement Guide to Reducing Carbon Emissions in Building and Construction*, 2021; *Whole-of-Life Embodied Carbon Assessment*, 2022).

The carbon literature supports the A1–A5 boundary used in this thesis because the comparison is between alternative CFS framing configurations at product, transport, and construction-stage levels. A wider whole-building LCA boundary would introduce envelope, services, operational energy, replacement, end-of-life, and beyond-boundary effects that are outside the framing configuration decision. The A1–A5 framing boundary is therefore appropriate for comparing steel mass, transport, installation, waste, and prefabrication-related effects, while remaining clearly separate from certified whole-building LCA.

The carbon literature therefore supports an explicit A1–A5 CFS framing boundary because the study compares alternative framing configurations, not whole-building life-cycle performance.

2.6: Construction Productivity, Time, and Process-Hour Modelling

The CFS framing process-hours can be estimated at various modelling levels of detail. Traditional project scheduling techniques like the Critical Path Method (CPM) view construction as a sequence of activities that depend on each other and can be helpful in estimating calendar duration, float, and critical activities (Liu et al., 2015). With BIM-based scheduling, this logic is pushed to another level, by connecting construction activities to the model objects – known as 4D BIM – which enables construction sequences and logistics to be visualised prior to construction. In the words of Doukari et al., 4D BIM consists of connecting the schedule tasks with the 3D model, which is in object oriented format, to enhance the logistical decision making and project delivery (Doukari et al., 2022). This study does not aim to create a complete CPM or 4D calendar programme. Rather, it employs a process-hour approach: time is the total production and installation time required for the CFS framing tasks (Doukari et al., 2022; Lam et al., 2001). This is appropriate since the research question is about the configuration's impact on the effort required to frame the CFS, not the entire project.

The number of components, number of repetitions of component families, setup requirements and complexity of fabrication detail affect factory time in panelised CFS construction. The panelised workflow includes the following steps: digital scheduling, roll forming, cutting, punching, labelling, assembly, quality checking, packaging and preparing for transport (Pan et al., 2008; Zolghadr et al., 2022). In prefabricated construction management literature, a few important management issues, which are not secondary details such as production planning, standardisation, logistics, and coordination, are repeatedly mentioned. Li et al.'s review of prefabricated construction management revealed prefabricated construction systems' research related to production planning, management methods, and implementation challenges (Li et al., 2014). In the current research , this allows modelling factory time as a function of the number of panels, number of trusses, burden of set up, number of repetition and opening/detailing complexity, rather than as a fixed duration.

Installation/ unloading/ positioning/ fixing/ connecting/ fit up and coordination activities are controlled by site time. CFS framing productivity is not only related to the number of panels and trusses at the site, but it is also impacted by the ability to find members, line up panels, join up panels, install roof trusses, and handle non-standard openings or penetrations (Bataglin et al., 2020). Being able to evaluate the decision to move construction off-site in a holistic manner has been identified through off-site construction research as extending beyond just direct cost, and also include benefits and constraints of the works from a labour, quality, logistics, safety, coordination, and project delivery perspective (Pan et al., 2008). Pan et al also demonstrate that the implementation barriers and practical conditions on the projects significantly affect the use of off-site construction methods by housebuilders, which reinforces the importance of considering the logistics and site-handling impacts where prefabrication is not automatically seen as reducing the duration of all construction activities (Pan et al., 2008).

Prefabrication can reduce site exposure and site labour demand, but the magnitude of this benefit is project dependent. The relevance of prefabrication to building projects is also evident in that it can reduce the construction waste, as shown by Jaillon (Jaillon et al., 2009) and this is also applicable to time modelling, as reduced rework, cutting and site waste handling can reduce the effort in the building stage (Lu et al., 2021). Productivity effects can also be brought about by repetition; in construction, research on the learning curve illustrates how a repetitive activity can affect labour productivity over time, particularly in cases where similar operations are performed many times. Lam et al. discussed learning and forgetting of repetitive construction operations and the effect of learning and forgetting on project productivity. Weather is also an actual cause of construction delay and productivity loss for site-based work. Schuldt et al (Lam et al., 2001; Schuldt et al., 2021). conducted a systematic review and determined that precipitation, extreme temperatures and high winds were significant weather-related factors that cause construction delays. The time model thus considers one or more weather downtime scenarios as an allowable variable to be used as a scenario term, not as an actual weather delay term (Ballesteros-Pérez et al., 2017; Schuldt et al., 2021).

Published CFS-specific installation productivity benchmarks are limited, particularly for panelised residential framing. For this reason, the time model in this study does not claim to be a universal CFS productivity estimator. Instead, it uses project-derived process-hour assumptions calibrated to the S2 case and supported by broader evidence from modular manufacturing, New Zealand off-site construction, and panelised timber installation studies (Lu et al., 2023). support the use of learning effects in modular construction manufacturing, showing that accumulated production experience can improve productivity, although learning rates may fluctuate. Shahzad's New Zealand off-site manufacturing research argues that off-site manufacturing can improve construction productivity in New Zealand, while also noting that adoption is constrained by industry barriers (Shahzad, 2011). Forsythe and Sepasgozar's study of prefabricated timber panel installation provides comparable panelised-construction evidence for measuring installation productivity, but it is not treated as direct CFS evidence (Forsythe & Sepasgozar, 2019). Therefore, the process-hour model is interpreted as a bounded comparative model for the case-study CFS framing system, not as a validated industry-wide productivity equation.

The time boundary in this study is therefore defined as CFS framing construction process-hours, not calendar days. The model includes factory production, setup, delivery/logistics, site installation, mobilisation, and weather allowance where applicable. It excludes full project scheduling dependencies such as subcontractor sequencing, procurement lead times, inspection delays, design approval time, crew availability, holidays, and overlapping activities outside the CFS framing scope. This distinction is important because two alternatives with similar process-hours may still produce different calendar durations depending on crew size, sequencing, parallel work, or site constraints. The objective of time should therefore be interpreted as a comparative measure of CFS framing production and installation effort relative to the S2 reference case, not as a complete construction programme.

Time objective:

$$T_{CFS}= T_{factory} + T_{setup} + T_{delivery/logistics} + T_{site\ installation} + T_{mobilisation} + T_{weather\ allowance}$$

Equation 2-5

where T_{CFS} is measured in process-hour

Table 2-6 Time-model literature review and relevance to CFS framing process-hours.

Study / source	Time-related focus	Method / concept	Relevance to this study	Limitation relative to this study
Kelley and Walker / CPM scheduling literature	Project sequencing and calendar duration	Critical Path Method, activity dependencies, float, project duration	Establishes that CFS framing process-hours can be modelled through activity networks and dependencies	this study does not create a full CPM calendar schedule
Eastman et al. / BIM Handbook	BIM-based planning and scheduling	4D BIM, linking model elements to construction activities	Supports linking building components to construction activities and sequencing logic	this study does not implement full BIM-based scheduling
Doukari et al. (2022)	Construction schedules in 4D BIM	Conventional and automated schedule creation	Supports the relationship between digital models, activity planning, and logistics	Not specific to CFS framing or multi-objective optimisation
Blismas et al. (2006)	Off-site production	Holistic evaluation of off-site production	Supports the argument that prefabrication affects labour, logistics, quality, and time	Does not provide CFS panel/truss

Study / source	Time-related focus	Method / concept	Relevance to this study	Limitation relative to this study
	benefits and constraints			process-hour equations
Pan et al. (2008)	Off-site construction uses by housebuilders	Adoption drivers and barriers	Supports the idea that off-site time benefits depend on practical implementation conditions	Not CFS-specific and not an optimisation model
Li et al. (2014)	Prefabricated construction management	Production planning, coordination, and implementation issues	Supports factory/site production-planning logic for prefabricated systems	Does not model panelised CFS process-hours
Jaillon et al. (2009)	Prefabrication and site waste	Survey and case-study evidence	Supports the link between prefabrication, reduced site work, and reduced waste/rework potential	Focuses on waste, not full-time modelling
Lam et al. (2001)	Repetitive construction productivity	Learning–forgetting effects	Supports the idea that repetition can influence productivity and activity duration	Not specific to CFS factory production.
Schuldt et al. (2021)	Weather-related construction delays	Systematic review of weather impacts	Supports treating weather exposure as a site-time risk or allowance	Does not predict project-specific weather delay
This study	CFS framing production and installation effort	Parametric process-hour model	Links R_p, R_t, OC, and P_f to CFS framing process-hours.	Does not estimate full project calendar duration

Note: this study uses process-hours rather than calendar duration. Process-hours represent factory production, setup, delivery/logistics, site installation, mobilisation, and weather allowance within the CFS framing scope. They do not represent a full CPM schedule or total project completion duration.

Direct CFS panel-installation productivity data are limited in the published literature, particularly for New Zealand residential projects using panelised wall frames and roof trusses. For this reason, the time objective is not presented as a universal CFS productivity estimator. Instead, it is a comparative process-hour response model calibrated to the S2 case-study boundary and supported by broader off-site, modular, DfMA, repetitive-construction, and weather-delay literature. This distinction is important because the model compares alternatives under a consistent set of process assumptions; it does not estimate full project calendar duration, subcontractor sequencing, procurement lead time, inspection delay, or detailed crew-resource scheduling.

Productivity and construction-time literature supports the use of process-hours as a comparative effort metric for this thesis. Process-hours are more appropriate than calendar duration because the model compares CFS framing configuration alternatives rather than full project programs with weather, crew sequencing, subcontractor interfaces, resource levelling, and non-CFS trade dependencies. The time objective is therefore interpreted as factory and site process effort within the framing boundary, not as total construction duration.

The productivity literature therefore supports process-hours as a comparative CFS framing effort measure, while keeping full calendar scheduling outside the model boundary.

2.7: Multi-Objective Optimisation and NSGA-II

Multi-objective optimisation is relevant to this thesis because early-stage CFS framing configuration decisions cannot be assessed reliably through a single objective. A configuration that reduces CFS framing process cost may not minimise A1–A5 embodied carbon or process-hours, while a configuration selected for low carbon may introduce factory, logistics, material, or site-installation penalties. Pareto-based optimisation is therefore appropriate because it generates a set of nondominated alternatives rather than forcing a single weighted objective before the trade-off structure is known (Coello et al., 2007; Deb, 2001).

NSGA-II is one of the most widely used Pareto-based evolutionary algorithms for engineering optimisation problems with competing objectives. It is commonly applied where the objective functions are nonlinear, interacting, and difficult to reduce to a closed-form optimum. Its role is to search the bounded CFS framing response space defined by (R_p), (R_t), (OC), and (P_f), and to generate a nondominated archive for engineering interpretation (Deb et al., 2002).

The use of NSGA-II does not imply that evolutionary optimisation is the only possible search approach for this four-variable response model. Full enumeration or dense grid search can be useful where variables are fully discrete, and the design space is small enough to evaluate exhaustively. However, the implemented response model includes continuous rationalisation and prefabrication variables, interaction terms, feasibility screening, and post-processing of representative alternatives. NSGA-II is therefore used as a defensible Pareto-search method, while the reported archive is assessed later through convergence diagnostics, grid-reference comparison, multi-seed stability checking, and S2-relative interpretation rather than assumed to be globally exact.

Other multi-objective algorithms are available, but NSGA-II is a defensible choice for this bounded three-objective problem. NSGA-III is designed for many-objective optimisation and uses reference points to guide diversity, which adds unnecessary reference-point tuning for the present three-objective model (Qingfu Zhang & Hui Li, 2007). MOEA/D decomposes the problem into scalar sub-problems and is powerful for many continuous optimisation contexts, but it requires a decomposition/weight-vector structure that is less intuitive for the discrete panel and truss scheme variables used here (Zitzler et al., 2001). SPEA2 uses strength Pareto fitness and an external archive, while multi-objective particle swarm optimisation extends swarm logic to multiple objectives but is naturally more suited to continuous decision spaces. Therefore, this study does not claim that NSGA-II is universally superior; it claims only that NSGA-II is appropriate for a three-objective, continuous bounded CFS framing response model in which the decision variables represent rationalisation intensity, detailing complexity, and prefabrication level (Coello et al., 2004; Deb & Jain, 2014).

Table 2-7 Multi-objective optimisation methods considered for the CFS framing screening framework.

Method	Key reference	Main principle	Strength	Limitation for this study
NSGA-II	Deb et al. (2002)	Pareto dominance, elitism, crowding distance	Well-established for two- and three-objective engineering optimisation; transparent and suitable for bounded continuous variables and post-processed reporting categories	Stochastic method; requires convergence and stability checks
NSGA-III	Deb and Jain (2014)	Reference-point-based many-objective optimisation	Effective for many-objective problems	Additional reference-point design is unnecessary for the present three-objective problem
MOEA/D	Zhang and Li (2007)	Decomposition into scalar sub-problems	Efficient for many continuous multi-objective problems	Requires decomposition/weight-vector structure; requires decomposition/weight-vector structure that is less intuitive for the decision-support aim
SPEA2	Zitzler et al. (2001)	Strength Pareto fitness and archive-based density estimation	Strong diversity preservation	External archive and density handling add complexity without clear benefit for this model
MOPSO	Coello et al. (2004)	Swarm-based multi-objective search	Useful for continuous decision spaces	less transparent than NSGA-II for the research workflow and less aligned with the existing implementation
DEAP	Fortin et al. (2012)	Python evolutionary-computation framework	Enables transparent implementation of custom evolutionary operators	Framework only; not an optimisation algorithm by itself
K-means clustering	Jain (2010)	Groups similar solutions in normalised objective space	Converts a large Pareto set into interpretable archetypes	Sensitive to scaling and chosen number of clusters
TOPSIS	Hwang and Yoon (1981); Behzadian et al. (2012)	Ranks alternatives by distance from ideal and anti-ideal points	Converts Pareto alternatives into stakeholder-specific rankings	Ranking depends on weights and normalisation; does not prove universal best solution

Table 2-7 shows that NSGA-II was selected as a defensible search method rather than as a claimed algorithmic contribution. The table also clarifies that full enumeration or grid-reference comparison can support diagnostic checking, but the study uses NSGA-II as the primary method because the implemented response model contains continuous variables, interaction effects, feasibility screening, and representative-solution post-processing

NSGA-II is selected because the research involves three objectives and a continuous bounded design vector, with derived reporting categories used only for interpretation the current optimisation uses a continuous design space K-means clustering and TOPSIS are used only as post-processing decision-support methods; they do not replace Pareto optimisation.

Recent construction and building-engineering studies confirm the practical relevance of NSGA-II and related multi-objective optimisation methods for problems involving cost, time, carbon, and uncertainty. (Wang et al., 2024) used NSGA-II for prefabricated component combination optimisation with cost, duration, and carbon-emission objectives. (Chaturvedi et al., 2020) applied NSGA-II to building design optimisation under operational uncertainties. (H. Song et al., 2024; Xu et al., 2025; Yuan et al., 2024) similarly demonstrates that prefabricated-building optimisation increasingly requires simultaneous treatment of CFS framing process-hours, cost, quality, and carbon emissions. These studies support the methodological choice of NSGA-II for the present research.

Existing literature already demonstrates the use of multi-objective optimisation, including NSGA-II, for cost, duration, and carbon trade-offs in prefabricated construction. The contribution of this thesis is therefore not the invention of a new optimisation algorithm. The contribution lies in applying a transparent Pareto-search workflow to a project-derived, New Zealand-context, prefabricated CFS residential framing problem using the CFS-specific variables (R_p), (R_t), (OC), and (P_f), followed by S2-relative classification, sensitivity analysis, Monte Carlo robustness testing, TOPSIS-based stakeholder ranking, and dashboard-supported interpretation.

The detailed implementation of the LHS-initialised NSGA-II workflow, including population settings, objective evaluation, convergence metrics, grid-reference comparison, multi-seed stability checking, and archive extraction, is reported in Section 3.12.

2.8: Latin Hypercube Sampling and Design-Space Coverage

Latin Hypercube Sampling (LHS) is relevant to this thesis as an initial design-space coverage method for NSGA-II. Its purpose is to reduce the risk that the generation-zero population is clustered in only one part of the four-variable design space. LHS is therefore used as a stratified initialisation method, not as an optimiser and not as proof of global optimality. (Lin & Tang, 2022; McKay et al., 1979).

Random initialisation can be appropriate in evolutionary algorithms, but it may produce uneven generation-zero coverage in bounded design spaces. LHS was selected in this study as a practical stratified initialisation method to improve marginal coverage of the four continuous variables at generation zero. This choice supports initial design-space coverage, but it is not treated as proof that LHS outperforms random initialisation for all runs (Iman & Conover, 1982; McKay et al., 1979; Morris & Mitchell, 1995).

LHS was introduced by McKay, Beckman and Conover as a stratified method that divides each variable range into equal-probability intervals and samples once from each interval. This gives more uniform marginal coverage than simple random sampling and makes LHS particularly suitable for computer experiments and evolutionary initialisation. The method is widely used because it is simple to implement, inexpensive to compute, and effective for ensuring that all parts of each variable range are represented in the initial sample (Iman & Conover, 1982; McKay et al., 1979; Morris & Mitchell, 1995).

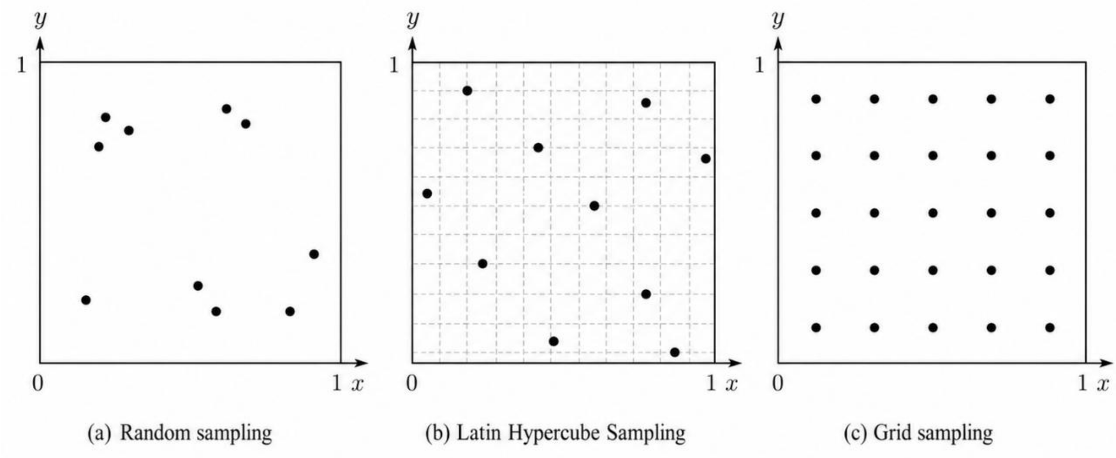

Figure 2-4 Conceptual comparison of random sampling, Latin Hypercube Sampling (LHS), and grid sampling in a two-dimensional design space (Source: Author's own schematic, based on McKay et al. (1979), Stein (1987), and Morris and Mitchell (1995))

illustrates the conceptual difference between random, LHS, and grid-based sampling in a two-dimensional design space. Random sampling may leave empty regions or generate clustering, while LHS provides better stratified coverage of each variable range. Grid sampling is highly regular but becomes inefficient as dimensionality increases and is less convenient for mixed discrete–continuous optimisation problems.

This mapping allows a single LHS design matrix to initialise the continuous four-variable optimisation problem. The scheme labels used later in the research are not sampled independently; they are derived from $R_p$and $R_t$after the rationalisation variables have been generated. Without requiring a separate sampling method for each variable type. The result is a generation-zero population that provides stratified coverage across R_p, R_t, OC, and P_f while still spanning the continuous range of opening complexity and prefabrication level. That is important here because the optimisation problem is not purely continuous and cannot rely on a standard single-variable sampling scheme (Lin & Tang, 2022; Navid et al., 2018; C. Song & Kawai, 2023).

The use of LHS also supports the applied screening purpose of the thesis. Since the framework is intended to compare early-stage CFS framing configurations rather than prove a closed-form mathematical optimum, the initial population should sample the feasible design domain broadly enough to reduce avoidable starting-point bias. LHS provides this initial spread while leaving Pareto search, feasibility filtering, convergence assessment, and archive interpretation to the NSGA-II workflow (Deb, 2004; Hamdan & Qudah, 2015; Helton & Davis, 2003; Hwang & Yoon, 1981; Morris & Mitchell, 1995).

Table 2-8 Comparison of LHS, random, and grid initialisation methods.

Method	Relevance to this thesis	Treatment in this study
Random initialisation	Simple baseline population-generation method	Not used as the primary initialisation method because it can cluster unevenly in a bounded four-variable space
Grid or dense enumeration	Useful diagnostic when variables are discretised	Used only as a reference/checking concept, not as the primary optimiser
Latin Hypercube Sampling	Stratified design-space coverage for continuous bounded variables	Used to initialise the NSGA-II generation-zero population
Sobol or quasi-random sampling	Alternative space-filling method	Not adopted because LHS provided a transparent and sufficient initialisation method for the implemented workflow

Table 2-8 clarifies that LHS is used as a practical initialisation choice rather than as a contribution in itself. The final optimisation results are not defended solely on the basis of LHS coverage. They are later assessed through convergence behaviour, grid-reference comparison, multi-seed stability, feasibility screening, and S2-relative classification.

The reviewed sampling literature supports LHS as a practical stratified sampling method for improving marginal coverage of a bounded design space compared with unstructured random sampling. However, LHS itself does not prove optimality, does not validate the physical accuracy of model coefficients, and does not guarantee discovery of the true global Pareto front. Existing LHS literature therefore supports the initialisation strategy but not the full credibility of the optimisation result. This study responds by using LHS only as a generation-zero design-space coverage method, while final optimisation quality is assessed separately through hypervolume convergence, grid-reference comparison, solution-count audit, and multi-seed stability checks.

2.9: Uncertainty Analysis and Monte Carlo

Monte Carlo simulation is relevant to this thesis because several inputs to CFS framing cost, carbon, process-hour, logistics, waste, and prefabrication response terms are uncertain rather than directly measured project records. Deterministic optimisation can identify nondominated alternatives under one coefficient set, but it can also give a false impression of precision if uncertain cost rates, carbon factors, productivity assumptions, transport conditions, and interaction coefficients are treated as fixed values. Robustness analysis is therefore needed to test whether selected Pareto alternatives remain defensible when plausible coefficient variation is considered (Helton & Davis, 2003; Montgomery & Douglas C., 2017).

In this study, Monte Carlo simulation is used after optimisation rather than as part of the optimiser. The NSGA-II procedure generates the nondominated archive, while Monte Carlo simulation tests selected representative solutions under bounded coefficient uncertainty. This distinction is important because the Monte Carlo stage does not create the Pareto front, does not recalibrate the model, and does not prove that the model physically replicates real project cost, carbon, or productivity outcomes. Where Latin Hypercube Sampling is used for uncertainty sampling or initial population coverage, it is treated as an established stratified sampling method rather than as a novel contribution (Hopfe & Hensen, 2011; McKay et al., 1979)

For early-stage CFS framing decision support, the main value of Monte Carlo simulation is comparative robustness. A deterministic low-cost or low-carbon solution may not be the most defensible practical option if its performance is sensitive to coefficient uncertainty. Conversely, a slightly less extreme solution may be more useful if it has a higher probability of remaining better than the S2 reference under the same uncertainty assumptions. This is particularly relevant for the present thesis because embodied-carbon and prefabricated-construction assessments are known to be sensitive to system boundary, emission factors, transport assumptions, construction-stage processes, and project-specific modelling choices (AbouHamad & Abu-Hamd, 2019; Hopfe & Hensen, 2011).

Table 2-9 Literature-supported uncertainty categories considered for robustness analysis.

Uncertainty category	**Example uncertain inputs**	**Objective affected**	**Monte Carlo role**	**Interpretation**
Cost rates	steel cost, factory processing rate, site labour rate, transport rate	Cost	Sample cost coefficients across plausible ranges	Tests whether cost ranking remains stable
Carbon factors	steel A1–A3 factor, freight factor, A5 installation factor	Carbon	Sample emission coefficients across plausible ranges	Tests sensitivity of carbon performance
Productivity assumptions	factory productivity, site installation productivity, prefabrication time-saving factor	Time, cost	Sample productivity-related coefficients	Tests whether process-hour savings remain robust
Logistics assumptions	transport distance, delivery/logistics time, mobilisation allowance	Cost, carbon, time	Sample project-specific logistics assumptions	Test's location and coordination sensitivity
Waste assumptions	steel waste rate, waste handling factor	Cost, carbon	Sample waste-related values	Tests influence of offcuts and waste allowance
Baseline comparison	S2 objective values under same uncertain inputs	All objectives	Compare candidate solutions against uncertain S2	Estimates probability of outperforming S2

Spearman and PRCC results are interpreted as screening-level indicators of monotonic association, not as proof of causal physical behaviour. This caution is necessary because the response model includes nonlinear terms, lower bounds, saturation effects, and interaction penalties. PRCC is therefore not used to infer direct sensitivity for derived panel_scheme or truss_scheme labels, because these are ordinal reporting categories rather than continuous optimiser-controlled variables. Grouped scheme-effect comparisons, OAT sensitivity, and engineering interpretation are used alongside rank-correlation results.

Although TOPSIS is widely used for multi-criteria decision support, it has known limitations. Rankings can be sensitive to the selected normalisation method, stakeholder weights, and the set of alternatives included, including the possibility of rank reversal when alternatives are added or removed. For this reason, TOPSIS is not used in this study as an optimiser or as proof of a universally best solution. It is used only as a transparent post-processing method for ranking a bounded set of representative Pareto alternatives under explicitly stated stakeholder-weight profiles.

The reviewed uncertainty literature supports Monte Carlo simulation, Spearman correlation, and PRCC as practical methods for propagating coefficient uncertainty and screening influential variables. However, these methods do not empirically validate the model; they only show how model outputs change when assumed input ranges are varied. This distinction is important because the CFS framing response model contains calibrated cost, carbon, productivity, logistics, waste, and prefabrication-response coefficients rather than measured project records for every input. This study responds by using Monte Carlo analysis as a post-optimisation robustness test for selected Pareto solutions relative to S2, while acknowledging that stronger validation would require measured factory, site, supplier, transport, and project-specific carbon data from additional CFS projects.

The Monte Carlo analysis in this study is therefore interpreted as bounded coefficient-uncertainty propagation, not as empirical validation. Measured cost records, product-specific verified carbon data, fabrication productivity records, and site installation records from multiple CFS projects would be required for external validation. The detailed Monte Carlo design, including coefficient ranges, number of simulations, paired comparison with S2, percentile reporting, coefficient of variation, and probability of outperforming S2, is defined in Section 3.15.

2.10: Dashboard and Visual Decision Support

Visual analytics literature emphasises that complex datasets are more useful when users can filter, select, compare, coordinate views, and move between overview and detail. Multi-objective results require users to compare alternatives, identify trade-off regions, inspect representative solutions, apply stakeholder priorities, and understand whether selected configurations are better or worse than a reference case. In this study, the dashboard is therefore positioned as an interpretation layer for model outputs rather than as an optimisation method (Gadelhak et al., 2017).

Visual analytics literature emphasises that complex datasets are more useful when users can filter, select, compare, coordinate views, and move between overview and detail. These ideas are directly relevant to Pareto-based CFS framing decision support because the user needs to inspect both the whole archive and selected alternatives. A practitioner should be able to see the overall cost–carbon–process-hour trade-off, then examine specific configurations such as cost-minimum, carbon-minimum, time-minimum, balanced, or robustness-preferred solutions (Few, 2006; Heer & Shneiderman, 2012; Shneiderman, 1996).

If multiple views are needed for multi-objective decision support, multiple views are linked together on the dashboard, not individual plots. Useful interface for this study should contain Pareto scatter plots, parallel-coordinate plots, representative-solution tables, S2-relative comparison cards, Monte Carlo uncertainty bands and ranking outputs. The framework of Munzner is relevant because it connects visualisation design with the data intention, the need for a user to do something with the data and how this is facilitated through the visualisation of the data and interactions with it (Munzner, 2014). This logic is key in the present dashboard as they might be asking different questions: a developer might ask which solution is the least expensive; a client might ask which solution has the lowest CFS framing A1–A5 upfront carbon; and a contractor might ask which solution has the lowest amount of hours in the process or exposure on site. The dashboard should, therefore, also allow for comparison and filtering, and not just present the final Pareto front in a static manner. The dashboard should allow users to inspect continuous R_p and R_t values while also filtering by the derived panel_scheme and truss_scheme categories for practical interpretation.

Visual Studio Code was used as the development environment for writing and managing Python optimisation and dashboard scripts. However, VS Code is not treated as the dashboard platform. The dashboard contribution is defined by the data pipeline, interface functionality, filtering, ranking, S2 comparison, and uncertainty visualisation, rather than by response model editor used to develop it.

The dashboard is not treated as a methodological novelty in itself. The contribution lies in connecting Pareto exploration, S2-relative comparison, component-level interpretation, sensitivity outputs, Monte Carlo robustness, TOPSIS-based ranking, and decision communication into one CFS framing-specific workflow. The dashboard implementation and exported model-output structure are defined later in the methodology and discussion chapters; the literature review establishes only why visual decision support is needed for interpreting multi-objective CFS framing results.

Table 2-10 Dashboard feature requirements from literature.

Requirement	Literature basis	Use in this study
Overview of alternatives	Visual analytics requires users to see the overall solution space before selecting details	Display the nondominated archive and S2 reference context
Filtering and selection	Users need to narrow complex datasets to relevant cases	Filter alternatives by objectives, $(R_p), (R_t), (OC), (P_f)$, and representative type
Details-on-demand	Users need detailed values after identifying a candidate solution	Show objective values, design-variable settings, component breakdown, and S2-relative change
Coordinated comparison	Multi-objective decisions require side-by-side comparison	Compare cost-minimum, carbon-minimum, time-minimum, balanced, and stakeholder-ranked alternatives
Robustness interpretation	Decision support must communicate uncertainty, not only deterministic values	Display Monte Carlo robustness, sensitivity results, and probability better than S2
Transparent boundary communication	Users must understand what the tool does and does not represent	State that outputs are early-stage screening results, not certified structural design, tender pricing, or whole-building LCA

Table 2-10 defines the dashboard requirements at the literature-review level. It does not describe the final dashboard interface. The implemented dashboard structure, exported data files, and user workflow are reported later, after the optimisation and robustness methods have been defined.

The dashboard literature supports visual decision support, but the thesis contribution lies in connecting Pareto alternatives, S2 comparison, robustness outputs, stakeholder ranking, and decision communication into a CFS framing-specific early-stage screening workflow.

2.11: Structured Gap Analysis

The reviewed literature supports the individual components required for this study: CFS structural behaviour, panelised/off-site construction, cost modelling, embodied-carbon estimation, construction productivity, NSGA-II optimisation, LHS sampling, Monte Carlo uncertainty analysis, and visual decision support. However, the reviewed studies do not provide a directly comparable framework that integrates these components for a project-derived CFS residential framing configuration problem in New Zealand. The gap is therefore defined as an integration and application gap, not as the absence of prior work in CFS, prefabrication, carbon assessment, optimisation, or dashboards individually.

The prefabrication and DfMA literature support the treatment of standardisation, repetition, production planning, and factory/site coordination as meaningful design decisions. These studies show that prefabrication can influence labour, productivity, logistics, waste, cost, and quality, but they also show that prefabrication benefits are project-dependent rather than automatic. This is important for the research because the model does not assume that higher prefabrication is always superior. Instead, prefabrication level, panel rationalisation intensity, truss rationalisation intensity, and opening/detailing complexity are treated as variables that may improve one objective while worsening another.

The reviewed cost, carbon, and time literature also supports the three objective functions used in this study. Cost studies justify separating material, factory processing, labour, logistics, transport, waste, and overhead rather than using a single whole-house cost rate. Embodied-carbon standards and guidance justify using a clearly stated A1–A5 boundary for CFS framing, while excluding operational, maintenance, end-of-life, and module D effects. Time and productivity literature supports the use of process-hours as a measure of production and installation effort, while making clear that this is different from a full CPM schedule or calendar duration

The optimisation literature supports the use of a Pareto-based method because the research problem involves competing objectives. NSGA-II is appropriate because it generates a set of non-dominated alternatives rather than forcing cost, carbon, and time into a single pre-weighted score. Closely related prefabricated-construction studies have used multi-objective optimisation for cost, duration, and carbon emissions, including NSGA-II-based prefabricated component optimisation and recent cost–duration–carbon optimisation frameworks. However, these studies are not specific to CFS residential panel rationalisation, truss rationalisation, opening/detailing, and prefabrication configuration. Similarly, LHS, Monte Carlo simulation, clustering, TOPSIS, and visual analytics are established methods, but they are usually applied as separate methodological tools rather than integrated into one CFS framing decision-support workflow.

Table 2-11 Structured gap matrix linking reviewed literature to the CFS framing optimisation framework.

Study	CFS/LGSF	CFS rationalisation variables	Cost	Carbon	Time	NSGAII/ MOO	LHS	Monte Carlo	Dashboard	Gap relative to research
Liang et al. (2022)	✓	✗	◐	◐	✗	◐	✗	✗	✗	CFS optimisation structural/member-level
Wang et al. (2024)	✗	◐	✓	✓	✓	✓	✗	✗	✗	Prefabricated components, but not CFS framing variables
Chaturvedi et al. (2020)	✗	✗	◐	◐	◐	✓	✗	◐	✗	Building design optimisation, not CFS production configuration
Yuan et al. (2024)	✗	◐	✓	✓	✓	✓/simulation	✗	✓	✗	Prefabricated process optimisation, not CFS panel/truss framework
Song et al. (2024)	✗	◐	✓	✓	✓	✓	✗	✗	✗	Assembled building optimisation, not CFS residential framing
Gadelhak et al. (2017)	✗	✗	◐	◐	◐	visualises MOO	✗	✗	✓	Dashboard support, not CFS optimisation model
this study	✓	✓	✓	✓	✓	✓	✓	✓	✓	Integrated CFS framing decision-support framework using R_p, R_t, OC, P_f, LHS-initialised NSGA-II, Monte Carlo robustness, and dashboard interpretation.

Legend: ✓ = directly addressed; ◐ = partially or indirectly addressed; – = not addressed / not a main focus.

Table 2-11 shows that the reviewed literature supports many individual components of the proposed framework, but that these components are usually treated separately. CFS studies provide the structural and material basis for light-gauge framing, but they are generally stronger in member-level or section-level optimisation than in early-stage production-configuration screening. Prefabricated-construction and DfMA studies address standardisation, production efficiency, logistics, site productivity, cost, time, and carbon, but they are usually not specific to project-derived panelised CFS residential framing. Similarly, NSGA-II, LHS, Monte Carlo simulation, TOPSIS, and dashboard-based visual decision support are established methods, but they are rarely integrated into one transparent workflow for CFS framing configuration decisions.

The gap addressed by this study is therefore an integration gap rather than the absence of any individual method. The reviewed literature establishes CFS structural design, off-site construction, DfMA, cost estimation, embodied-carbon assessment, construction productivity modelling, multi-objective optimisation, LHS sampling, Monte Carlo uncertainty analysis, TOPSIS ranking, and dashboard-based decision support as separate research areas. However, no directly comparable study was identified that integrates project-derived CFS residential framing data, the four configuration variables $(R_p), (R_t), (OC), (P_f)$, simultaneous CFS framing process cost, CFS framing A1–A5 embodied carbon, CFS framing process-hours, LHS-initialised NSGA-II optimisation, Monte Carlo robustness testing, representative Pareto-solution interpretation, stakeholder ranking, and dashboard-supported decision communication within one transparent New Zealand case-study framework.

This gap leads directly to the methodology in Chapter 3. The following chapter defines the fixed S2 case-study boundary, decision-vector formulation, objective functions, coefficient audit, normalisation convention, NSGA-II configuration, sensitivity methodology, Monte Carlo robustness design, TOPSIS ranking logic, and dashboard export logic used to implement the framework.

Chapter 3: Research Methodology and Response-Model Formulation

Chapter 3 defines the methodology used to convert the project-derived CFS residential framing package into a bounded cost–carbon–process-hour decision-support model. The chapter establishes the fixed S2 case-study boundary, the four-variable decision vector, the objective functions, the response coefficients, the normalisation convention, the feasibility-screening logic, the LHS-initialised NSGA-II optimisation workflow, the sensitivity and Monte Carlo robustness methods, and the dashboard-output requirements.

The purpose of this chapter is implementation traceability rather than literature review. Chapter 2 established the literature basis for CFS framing, prefabrication, cost, carbon, process-hours, optimisation, uncertainty analysis, ranking, and visual decision support. Chapter 3 defines how those concepts were operationalised in the implemented response model. Standard method theory is therefore not repeated except where needed to specify the actual modelling choices used in this thesis.

The methodology is interpreted as an early-stage screening framework for CFS framing configuration alternatives. It does not produce certified structural designs, contractor tender prices, whole-building life-cycle assessments, or full construction schedules. Instead, it provides a transparent, coefficient-audited, and uncertainty-tested basis for comparing CFS framing configurations within the fixed New Zealand case-study boundary.

3.1: Research Design

This study uses a deterministic multi-objective response model with post-optimisation sensitivity and robustness analysis. The deterministic component evaluates alternative CFS framing configurations and uses an LHS-initialised NSGA-II algorithm to generate a nondominated approximation of the cost–carbon–process-hour trade-off surface. The sensitivity and Monte Carlo stages are then used to interpret the stability, uncertainty, and practical robustness of selected Pareto alternatives rather than to replace the deterministic optimisation process.

The objective functions are parametric response models developed for comparative decision support. They are not first-principles structural mechanics equations, contractor tender-price equations, certified life-cycle assessment equations, or empirically validated productivity equations. Their purpose is to provide a transparent and reproducible basis for comparing CFS framing configuration alternatives under a consistent case-study boundary. The results are therefore interpreted as model-calculated comparative outputs within the stated assumptions and coefficient ranges.

Because the problem contains three competing objectives, the objectives are not combined into a single pre-weighted scalar function. Instead, NSGA-II is used to generate a nondominated approximation of the cost–carbon–process-hour trade-off surface through non-dominated sorting and crowding-distance preservation. This is based on the universally adopted NSGA-II formulation given by Deb et al. for elitist multi-objective evolutionary optimisation (Deb et al., 2002).

Initially the population is generated through Latin Hypercube Sampling and not from random sampling alone. It is used to enhance the initial (or starting) coverage of the 4-D design space before the evolutionary search process ensues. Appropriately, LHS is used here to move into each dimension of the input and minimize the chances that the initial population of the NSGA-II algorithm is concentrated in just one part of the design space. This is what was originally used LHS as a sampling method for computer code experiments and simulation studies (McKay et al., 1979).

The implemented framework defines four active optimiser variables: panel rationalisation intensity, truss rationalisation intensity, opening/detailing complexity, and prefabrication level. In response model, these are represented as R_p, R_t, OC, and P_f, with R_p and R_t treated as continuous rationalisation intensities and mapped to practical panel/truss scheme labels only for reporting and dashboard interpretation. Response model also defines S2 as a model-calculated reference case and explicitly states that S2 is used for post-optimisation comparison, not as a hidden calibration scaler.

The environmental boundary is for A1 – A5 framing related impacts. This means that the carbon objective designs incorporate the use-stage carbon of the product (if applicable), transport to site, the factory/site process (if applicable), construction-stage carbon, and the carbon of the waste allowances; however, the carbon of the use and end-of-life stages, as well as the operational carbon, is not incorporated. This conforms to the current product-transport-construction stage "upfront carbon" building carbon accounting system. The current framework for EN 15978, to assess building environmental performance continues to apply a life-cycle assessment (LCA) approach, and the NZGBC's methodology for assessing embodied carbon also classifies A1–A5 as upfront carbon(European Committee for Standardization (CEN), 2026; New Zealand Green Building Council, 2024).

The research design consists of five linked stages: case-study and S2 definition, deterministic response modelling, LHS-initialised NSGA-II optimisation, sensitivity and uncertainty analysis, and dashboard-based interpretation. This sequence allows every research result to be shown to be derived from a controlled input variable, via a mathematical equation, to a specific code component and output file.

Figure 3-1 Overall methodology workflow for the CFS framing optimisation and decision-support framework. The workflow begins with the fixed CFS case-study boundary and S2 reference scenario, then defines the design variables, assumptions, feasibility screening, interaction response model, and cost–carbon–time objective evaluation. The raw objective values are normalised against S2 and searched using LHS-initialised NSGA-II. The resulting nondominated archive is classified against S2, validated using convergence and stability metrics, tested through sensitivity and Monte Carlo robustness analysis, and exported for dashboard-based decision support.

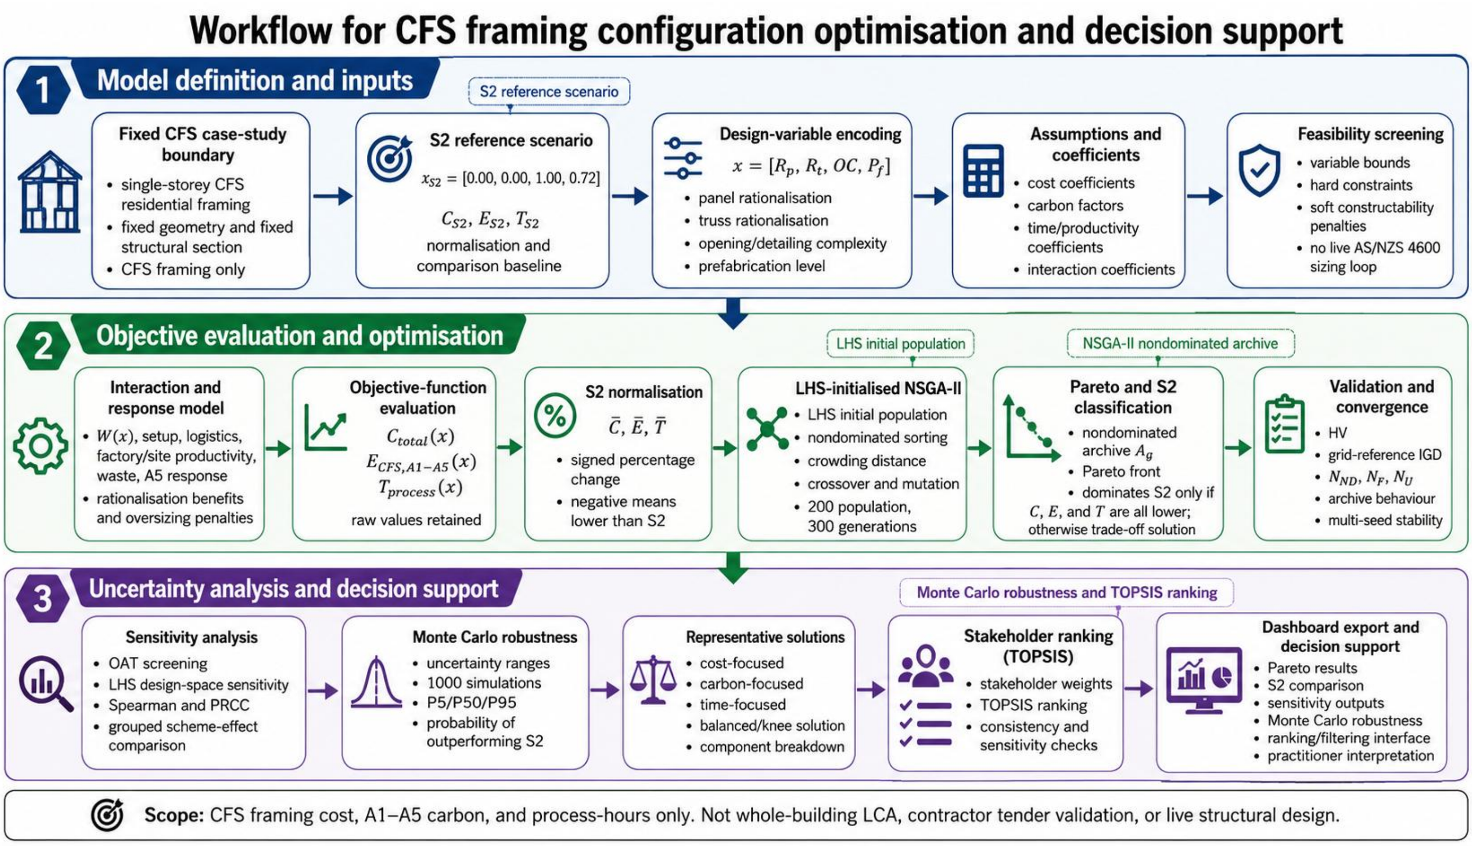

Figure 3-1 Methodology workflow for the CFS framing configuration-screening framework

The workflow shows how the fixed case-study boundary and model-calculated S2 reference are used to define the decision variables, response coefficients, feasibility checks, cost–carbon–process-hour objectives, LHS-initialised NSGA-II search, nondominated archive extraction, S2-relative classification, sensitivity and Monte Carlo robustness testing, TOPSIS ranking, and dashboard-supported interpretation. The figure represents the methodology sequence rather than software architecture.

Table 3-1 Method-to-output mapping for the research design.

Method stage	**Model role**	**Controlled / derived / fixed / excluded**	**Code/output trace**	**Research evidence**
Case-study boundary definition	Defines the fixed CFS framing scope	Fixed	Baseline parameter class / fixed boundary table	Section 3.2, S2 baseline table
Decision-vector encoding	Defines what NSGA-II directly controls	Controlled	Decision vector and variable bounds	Section 3.3
Rationalisation response mapping	Converts R_p and R_t into effective panel-class and truss-type response quantities	Derived	R_p and R_t bounds	Section 3.3 and 3.8
Parametric objective-response model	Calculates modelled cost, carbon, and process-hours for each candidate configuration	Derived response	Cost/carbon/time functions	Sections 3.5-3.7

Method stage	Model role	Controlled / derived / fixed / excluded	Code/output trace	Research evidence
Weight and interaction response model	Links rationalisation, prefabrication, opening/detailing complexity, logistics, and steel-mass response through audited coefficients	Derived response	Weight and interaction terms	Section 3.8
Feasibility screening	Applies variable bounds and fixed case-study constructability thresholds	Fixed thresholds + derived penalties	Feasibility/penalty report	Section 3.11
LHS initialisation	Improves generation-zero design-space coverage	Method	Initial population audit	Section 3.12
NSGA-II optimisation	Generates non-dominated alternatives	Method	Pareto front output	Chapter 5
Grid and convergence validation	Tests Pareto approximation quality	Method validation	HV, IGD, grid reference, multi-seed outputs	Chapters 5-6
Sensitivity analysis	Tests variable influence	Interpretation method	OAT, Spearman, PRCC, grouped effects	Chapter 6
Monte Carlo robustness	Tests coefficient uncertainty	Robustness method	P5/P50/P95, CV, probability better than S2	Chapter 6
Dashboard interpretation	Converts outputs into decision-support workflow	Decision-support output	Exported CSV/PNG/dashboard data	Chapter 7

Table 3-1. Method-to-output mapping for research design. The table links each methodological stage to its role in the model, its controlled/derived/fixed status, the corresponding implementation output, and the research section where the evidence is reported. The checks reported in this methodology are therefore interpreted as computational verification, practical reasonableness assessment, and bounded robustness testing rather than external empirical validation.

Table 3-2 Literature-to-model traceability matrix.

Model element	Literature support	Use in this study	Boundary
Fixed CFS framing system	AS/NZS 4600, CFS design literature	Establishes fixed structural boundary	No live member redesign
CFS configuration variables	Prefabrication and DfMA literature	$(R_p), (R_t), (OC), (P_f)$	Response variables, not direct fabrication commands
Cost objective	Construction cost/off-site literature	CFS framing process cost	Not whole-house tender price
Carbon objective	EN 15978 / ISO / NZGBC / steel LCA literature	CFS framing A1–A5 carbon	Not whole-building LCA

Model element	Literature support	Use in this study	Boundary
Process-hour objective	Productivity, CPM, 4D BIM, off-site literature	Factory/site process-hours	Not calendar programme
Optimisation and decision support	NSGA-II, LHS, Monte Carlo, visual analytics	Pareto + robustness + dashboard workflow	Computational/model-based validation

Table 3-2 links the reviewed literature domains to the implemented response-model components. It is placed in the methodology chapter because its function is to demonstrate model traceability rather than to extend the literature review. The table also clarifies the boundary of each component, including the distinction between CFS framing process cost and whole-house tender cost, A1–A5 framing carbon and whole-building LCA, process-hours and calendar duration, and dashboard-supported decision communication and certified engineering design.

3.2: Case-Study Boundary and Baseline S2

3.2.1: Purpose of the case-study boundary

The S2 configuration is used as the model-calculated reference case for normalisation, S2-relative comparison, and robustness testing. It is not treated as measured project performance, contractor tender cost, verified carbon certification, or observed site productivity. Its role is to provide a consistent reference configuration within the fixed case-study boundary. Later chapters therefore refer to S2 as the reference case without repeating this full boundary statement.

This study is applied to a fixed single-storey cold-formed steel residential framing case study. The optimisation boundary is limited to the CFS framing package, including wall-panel framing, roof-truss framing, framing steel mass, CFS framing process cost, CFS framing A1–A5 carbon, and CFS framing process-hours. This boundary is intentionally narrower than a whole-house assessment because the purpose of the model is to compare alternative CFS framing configurations rather than redesign the complete building.

The model does not optimise architectural layout, structural member sizing, connection design, bracing design, foundation design, services, linings, cladding, operational energy, end-of-life impacts, or full construction programme duration. These elements are held outside the optimisation boundary so that the model remains a transparent CFS framing decision-support framework rather than a whole-building design, pricing, or compliance tool.

The fixed case-study geometry is represented by the footprint area, A_f, and envelope/framing area, A_e:

$$A_f = 98.292 \text{ m}^2$$

$$A_e = 217.2555 \text{ m}^2$$

where A_f is the building footprint area and A_e is the CFS envelope/framing area used for process and carbon calculations.

The fixed framing quantities are the number of wall panels, number of roof trusses, and S2 reference steel mass:

$$N_p = 48$$

$$N_t = 58$$

$$W_{S2} = 2195.264 \text{ kg}$$

where N_p is the number of wall panels, N_t is the number of roof trusses, and W_{S2} is the reference CFS steel mass for the S2 case.

The fixed section and steel-mass boundary mean that the optimiser evaluates CFS framing configuration effects rather than structural member redesign. The model therefore compares rationalisation, detailing complexity, prefabrication level, logistics, and process-response effects within a fixed case-study framing system. It does not resize CFS members, redesign connections, verify bracing, or replace consent-level structural design.

Measured records for factory labour, site labour, transport cost, contractor pricing, and construction-stage carbon were not available. Therefore, the S2 cost, carbon, process-hours, and prefabrication-level values were calculated or assigned within the response model using documented coefficients and scenario assumptions. S2 is used as a consistent reference for normalisation, percentage-change reporting, Monte Carlo paired comparison, and dashboard interpretation; it is not treated as a measured contractor benchmark or universal New Zealand CFS baseline.

The value $P_{f,S2} = 0.72$ was selected as a response-model reference representing a panelised CFS framing package in which wall panels and roof truss components are substantially factory-prepared, while installation, alignment, fixing, bracing interfaces, service coordination, and site fit-up remain site-based activities. It is therefore not a measured percentage of total house construction completed off-site. The value is used as a calibrated reference point for the factory–site balance in the response model and should be recalibrated where measured fabricator production data are available.

3.2.2: Fixed case-study notation

The S2 configuration is used as the model-calculated reference scenario for normalisation and post-optimisation comparison. It represents the baseline CFS framing configuration against which alternative solutions are assessed. S2 is not treated as measured whole-house construction data, and it is not used as a hidden calibration factor inside the optimisation process.

The S2 decision vector is defined as:

$$x_{S2} = \left[R_{p,S2},\ R_{t,S2},\ OC_{S2},\ P_{f,S2}\right]$$

Equation 3-1

For the case-study reference condition:

$$x_{S2} = [0.00,\ 0.00,\ 1.00,\ 0.72]$$

where $R_{p,S2} = 0.00$is the S2 panel rationalisation intensity, $R_{t,S2} = 0.00$is the S2 truss rationalisation intensity, $OC_{S2} = 1.00$is the reference opening/detailing complexity, and $P_{f,S2} = 0.72$is the reference prefabrication level.

Response model-calculated S2 benchmark values are:

$$C_{S2} = C_{\text{total}}(x_{S2}) = 18335.47 \text{ NZD}$$

$$E_{S2} = E_{A1-A5}(x_{S2}) = 5366.79 \text{ kgCO}_2\text{e}$$

$$T_{S2} = T_{\text{process}}(x_{S2}) = 93.92 \text{ h}$$

These values are calculated directly from the model at the S2 reference vector. The cost value includes steel material, factory processing, site labour, factory setup, transport, delivery/logistics, site mobilisation, waste handling, and overhead. The carbon value includes A1–A3 steel, A4 transport, factory energy, A5 installation, and waste allowance. The time value includes factory production, setup, delivery/logistics, site installation, site mobilisation, and the modelled weather downtime allowance.

3.2.3: S2 Reference Breakdown

The S2 cost benchmark is calculated as:

$$C_{S2} = C_{\text{steel}} + C_{\text{fac}} + C_{\text{site}} + C_{\text{setup}} + C_{\text{trans}} + C_{\text{del}} + C_{\text{mob}} + C_{\text{waste}} + C_{\text{lift}} + C_{\text{oh}}$$

Equation 3-2

Table 3-3 Cost benchmark for S2.

Cost component	Symbol	Value	Unit
Steel material	C_{steel}	5,268.63	NZD
Factory processing	C_{fac}	6,822.88	NZD
Site labour	C_{site}	1,989.40	NZD
Factory setup	C_{setup}	1,495.00	NZD
Transport	C_{trans}	823.18	NZD
Delivery/logistics	C_{del}	98.08	NZD
Site mobilisation	C_{mob}	144.00	NZD
Waste handling	C_{waste}	27.44	NZD
Lifting equipment	C_{lift}	0.00	NZD
Overhead	C_{oh}	1,666.86	NZD
Total S2 cost	C_{S2}	18,335.47	NZD

Table 3-3 reports the model-calculated S2 cost benchmark used as the cost denominator for normalisation and S2-relative comparison. The values are taken from the final implemented cost-objective response at the S2 reference vector $x_{S2} = [0,0,1.00,0.72]$. The factory-processing, site-labour, transport, logistics, and overhead terms therefore reflect the final response-model structure rather than an earlier static baseline decomposition. The total S2 process cost is $C_{S2} = 18{,}335.47\text{NZD}$ and is used consistently in the optimisation results, percentage-change calculations, and dashboard comparison.

The S2 carbon benchmark is calculated as:

$$E_{S2} = E_{A1-A3} + E_{A4} + E_{\text{fac}} + E_{\text{site}} + E_{A5} + E_{\text{waste}}$$

Equation 3-3

Table 3-4 Carbon benchmark for S2.

Carbon component	Symbol	Value	Unit
A1–A3 steel	E_{A1-A3}	4610.05	$kgCO_{2e}$
A4 transport	E_{A4}	35.86	$kgCO_{2e}$
Factory energy	E_{fac}	14.36	$kgCO_{2e}$
Site energy	E_{site}	0	$kgCO_{2e}$
A5 installation	E_{A5}	695.22	$kgCO_{2e}$
Waste allowance	E_{waste}	230.5	$kgCO_{2e}$
Total S2 carbon	E_{S2}	5366.79	$kgCO_{2e}$

The S2 process-time benchmark is calculated as:

$$T_{S2} = T_{\text{factory}} + T_{\text{setup}} + T_{\text{delivery}} + T_{\text{site}} + T_{\text{mobilisation}} + T_{\text{weather}}$$

Equation 3-4

The time value represents CFS framing process-hours only. It includes factory production, setup, delivery/logistics, site installation, site mobilisation, and the modelled weather downtime allowance. It is not a critical-path construction programme and does not represent the total calendar duration of house construction.

Table 3-5 S2 Baseline quantities.

Parameter	Symbol	Value	Unit	Model role	Status
Footprint area	A_f	98.292	m²	Defines case-study building scale	Fixed
Envelope/framing area	A_e	217.2555	m²	Used in CFS framing process and A5 carbon calculations	Fixed
Wall panel count	N_p	48	count	Used for panel repetition and constructability interpretation	Fixed
Roof truss count	N_t	58	count	Used for truss repetition and constructability interpretation	Fixed
S2 steel mass	W_{S2}	2195.264	kg	Used for material cost and A1–A3 steel carbon	Fixed/reference
S2 panel rationalisation intensity	$R_{p,S2}$	0.00	index	Reference panel rationalisation condition	Reference
S2 truss rationalisation intensity	$R_{t,S2}$	0.00	index	Reference truss rationalisation condition	Reference
S2 opening/detailing complexity	OC_{S2}	1	dimensionless	Reference opening/detailing condition	Reference

Parameter	Symbol	Value	Unit	Model role	Status
S2 prefabrication level	$P_{f,S2}$	0.72	fraction	Reference prefabrication level	Assigned medium-prefabrication reference level within response model
S2 total cost	C_{S2}	18335.47	NZD	Cost normalisation and comparison denominator	Model-calculated reference
S2 total carbon	E_{S2}	5366.79	$kgCO_{2e}$	Carbon normalisation and comparison denominator	Model-calculated reference
S2 total process time	T_{S2}	93.92	h	Time normalisation and comparison denominator	Model-calculated reference

Note: S2 is a project-derived, model-calculated reference configuration. It is used for internal comparison and normalisation only; it is not an external measured benchmark.

3.2.4: Normalisation Role of S2

The S2 values are used as denominators for objective normalisation. The normalised reference condition is therefore:

$$\bar{C}(x_{S2}) = 1, \quad \bar{E}(x_{S2}) = 1, \quad \bar{T}(x_{S2}) = 1$$

For any candidate solution **x**, later sections define the normalised objectives as:

$$\bar{C}(x) = \frac{C_{\text{total}}(x)}{C_{S2}}$$

$$\bar{E}(x) = \frac{E_{A1-A5}(x)}{E_{S2}}$$

$$\bar{T}(x) = \frac{T_{\text{process}}(x)}{T_{S2}}$$

Equation 3-5

These equations are introduced here only to clarify the role of S2. The full normalisation and percentage-change formulation is defined later in Section 3.9.

3.3: Decision Vector and Variable Encoding

The optimisation framework represents each candidate CFS framing configuration using a compact decision vector that can be evaluated by the cost, embodied carbon, and construction-time response equations. The purpose of the decision vector is not to describe a complete architectural or structural redesign of the case-study dwelling. Instead, it defines the set of framing-configuration decisions that are varied within the fixed project boundary established in the preceding methodology sections.

The decision vector is defined as:

$$x = [R_p, R_t, OC, P_f]$$

Equation 3-6

where R_p is the panel rationalisation intensity, R_t is the truss rationalisation intensity, OC is the opening/detailing complexity index, and P_f is the prefabrication level. The variables R_p, R_t, OC, and P_f are treated as continuous bounded variables in the implemented optimisation framework.

The four variables in $(\mathbf{x})$ are treated as bounded response-model variables rather than informal judgement scores. (R_p) and (R_t) represent the rationalisation intensity applied to wall-panel and roof-truss families, respectively. Their purpose is to control the degree of standardisation and repetition represented in the response model, while the derived panel-scheme and truss-scheme labels are used only for reporting and interpretation. (OC) represents opening/detailing complexity relative to the S2 reference condition, with lower values representing simplified or more standardised detailing and higher values representing increased opening, trimming, lintel, jamb, reinforcement, and local fit-up burden. (P_f) represents the level of prefabrication and the associated shift of work between factory-controlled preparation and site installation.

This distinction is important because the optimisation does not directly redesign the architectural layout or structural member sizes. The variables are used to explore configuration effects within the fixed CFS framing package. The project fabrication drawings provide the physical basis for the wall-panel layout, opening/detailing conditions, and roof-truss diversity, while the response model defines how rationalisation, detailing complexity, and prefabrication level affect cost, A1–A5 carbon, and process-hours. The variables are therefore measurable within the model boundary, but they should not be interpreted as automatically extracted CAD quantities for every hypothetical design unless such extraction is separately implemented.

Table 3-6 Decision-variable interpretation and measurability within the response model.

Variable	Engineering interpretation	Project-derived basis	Model role	Reporting distinction
$(\boldsymbol{R_p})$	Panel rationalisation intensity	Wall-panel family repetition and standardisation logic	Controlled optimiser variable	Panel scheme is derived for reporting only
$(\boldsymbol{R_t})$	Truss rationalisation intensity	Roof-truss family repetition and span/type diversity	Controlled optimiser variable	Truss scheme is derived for reporting only
$(\boldsymbol{OC})$	Opening/detailing complexity	Openings, lintels, jambs, trimming, reinforcement, local fit-up burden	Controlled response-model variable relative to S2	Not a direct window-area ratio
$(\boldsymbol{P_f})$	Prefabrication level	Factory–site work split for the CFS framing package	Controlled optimiser variable	Not a calendar-schedule variable

Table 3-6 clarifies that R_p, R_t, OC, and P_f are the variables directly controlled by the optimiser. Panel scheme and truss scheme are derived interpretation labels and are not independent optimiser variables.

The design vector may also be expressed in component form as:

$$x = [x_1, x_2, x_3, x_4] = [R_p, R_t, OC, P_f]$$

Equation 3-7

The feasible design domain is defined as:

$$\Omega = \{ (R_p, R_t, OC, P_f) \mid 0.00 \leq R_p \leq 1.00,\ 0.00 \leq R_t \leq 1.00,\ 0.80 \leq OC \leq 1.30,\ 0.50 \leq P_f \leq 0.90 \}$$

Equation 3-8

The individual variable bounds are:

$$R_p \in [0.00, 1.00]$$

$$R_t \in [0.00, 1.00]$$

$$OC \in [0.80, 1.30]$$

$$P_f \in [0.50, 0.90]$$

Equation 3-9

The S2 reference configuration is defined as:

$$x_{S2} = [0.00, 0.00, 1.00, 0.72]$$

In this reference vector, $R_p = 0.00$ and $R_t = 0.00$ represent the S2-like panel and truss rationalisation condition, $OC = 1.00$ represents the reference opening/detailing complexity condition, and $P_f = 0.72$ represents the medium-prefabrication reference condition. The S2 vector is used for objective normalisation and post-optimisation comparison. It is not used as an additional hidden optimisation constraint.

The practical panel_scheme and truss_scheme labels derived from reporting categories only. They are not direct optimiser-controlled variables. The optimiser operates on the continuous rationalisation intensities $R_p$and R_t, while panel_scheme and truss_scheme are used later to support interpretation, plotting, tables, and dashboard filtering.

The panel rationalisation intensity R_p represents the degree of wall-panel standardisation in the CFS framing configuration. A value of $R_p = 0.00$ represents the S2-like reference condition, while higher values represent stronger panel rationalisation. Within the response model, R_p affects effective panel-class count, repetition, setup effort, factory production response, site installation response, material-weight response, waste response, and logistics interaction.

The truss rationalisation intensity R_t represents the degree of roof-truss span-family rationalisation. A value of $R_t = 0.00$ represents the S2-like reference condition, while higher values represent stronger truss-family rationalisation. Within the response model, R_t influences effective truss-type count, fabrication repetition, setup effort, installation effort, material-weight response, waste response, and logistics interaction.

3.3.1: Effective panel-class and truss-type quantities

The panel and truss rationalisation variables are translated into effective response quantities before being used in the cost, carbon, and process-hour equations. The effective number of panel classes is defined as a continuous response quantity: $N_{p,\text{class}}\left(R_p\right) = 4 - 2R_p$

where $N_{p,\text{class}}(R_p)$ represents the effective number of wall-panel classes associated with the selected panel rationalisation intensity. At $R_p = 0$, the model represents the S2-like condition with four effective panel classes. At $R_p = 1$, the model represents a highly rationalised condition with two effective panel classes. Intermediate values are interpreted as response-model quantities rather than literal fabrication schedules.

$$N_{t,\text{type}}(R_t) = 17 - 3R_t$$

where $N_{t,\text{type}}(R_t)$ represents the effective number of roof-truss span/type families associated with the selected truss rationalisation intensity. At $R_t = 0$, the model represents the S2-like condition with seventeen effective truss types. At $R_t = 1$, the model represents a more rationalised condition with fourteen effective truss types. These quantities are not direct fabrication commands and should not be read as requiring fractional panel classes or fractional truss types. They are continuous response-model variables used to represent the effect of rationalisation on repetition, setup effort, waste, material response, installation effort, and logistics interaction.

The opening/detailing complexity index OCrepresents the level of manufacturability and detailing complexity associated with openings, lintels, jamb studs, trimming members, service penetrations, and local coordination. It does not change the number of openings in the case study. Lower OCvalues represent greater standardisation of opening details, while higher OCvalues represent more customised detailing, local reinforcement, and coordination burden.

The opening/detailing complexity variable OC is a bounded design-for-manufacture proxy rather than an architectural opening ratio. The number of openings is fixed in the case-study boundary; OC represents the degree of standardisation or customisation in lintels, jamb studs, trimmers, service penetrations, CNC cutting patterns, local reinforcement, and site fit-up requirements. The range OC∈ [0.80,1.30] was selected to represent a practical detailing spectrum from highly standardised repeatable opening modules (OC=0.80) to heavily customised, non-repeated detailing with increased trimming, reinforcement, and coordination burden (OC=1.30). Therefore, OC=1.30 should not be interpreted as directly adding 30% to every objective; it is a bounded response-model input whose cost, carbon, and time effects are applied through the relevant opening/detailing coefficients.

The prefabrication level P_f represents the degree to which the CFS framing package is shifted toward off-site production. Higher P_f values can reduce site installation effort and site-stage carbon effects, but may also increase factory processing demand, logistics coordination, handling requirements, and factory-capacity pressure. This formulation therefore allows the optimisation model to represent both benefits and penalties rather than forcing all objectives to improve in the same direction. The prefabrication level P_f is bound between 0.50 and 0.90 to represent the practical range of factory-prepared CFS framing considered in this case-study model. The lower bound, $P_f = 0.50$, represents a low-prefabrication anchor in which a larger share of framing effort remains site-dependent. The S2 reference value, $P_f = 0.72$, represents the medium-prefabrication case-study condition, where wall panels and roof-truss elements are substantially factory-prepared but still require site installation, connection, alignment, bracing interface, and coordination activities. The upper bound, $P_f = 0.90$, represents a high-prefabrication anchor in which most CFS framing work is shifted toward factory-controlled production, while recognising that transport, handling, logistics, and site connection work remain necessary. Therefore, P_f is interpreted as a dimensionless response-model variable controlling the factory–site work balance, not as a directly measured percentage of total building construction completed off site.

Opening/detailing complexity, (OC), is treated in this thesis as a bounded response-model variable referenced to the S2 framing condition, where (OC=1.0) represents the adopted S2 opening/detailing burden. It is not used as an informal visual judgement score. The physical basis of (OC) is the additional fabrication and site-fit-up burden associated with openings, lintels, jamb studs, trimming members, local reinforcement, angle plates, non-standard panel conditions, and local connection/detailing requirements visible in the project-derived fabrication drawings and panel data.

For the reported optimisation run, (OC) is implemented as a controlled scenario variable rather than as a fully automated CAD-extracted index. This is an important boundary condition. The model tests how lower or higher opening/detailing burden would affect CFS framing process cost, A1–A5 carbon, and process-hours relative to S2, but it does not claim to automatically calculate (OC) from every CAD drawing. A fully automated implementation would require a separate extraction procedure based on measurable drawing quantities such as opening count, opening area, non-standard panel count, lintel and jamb frequency, trimmer count, cut-length diversity, reinforcement details, and local connection/detailing indicators.

Within this thesis, the defensible interpretation is therefore that (OC) is a calibrated response-model proxy anchored to observable opening/detailing mechanisms, not a universal CAD-derived metric. The OC-fixed diagnostic in Chapter 6 was included specifically to test whether the optimisation results depend only on reducing (OC) below the S2 reference value.

Table 3-7 Design variable encoding used in the optimisation framework.

Component	**Symbol**	**Type**	**Range / values**	**Unit**	**Model role**
Panel rationalisation intensity	R_p	Continuous optimiser variable	0.00–1.00	Dimensionless	Represents progressive wall-panel standardisation and repetition relative to S2
Truss rationalisation intensity	R_t	Continuous optimiser variable	0.00–1.00	Dimensionless	Represents progressive roof-truss span-family

Component	Symbol	Type	Range / values	Unit	Model role
					rationalisation relative to S2
Opening/detailing complexity	OC	Continuous optimiser variable	0.80–1.30	Dimensionless	DfMA/detailing complexity proxy for lintels, jambs, trimmers, service penetrations, cutting, reinforcement, and site fit-up
Prefabrication level	P_f	Continuous optimiser variable	0.50–0.90	Fraction	Represents the degree of off-site prefabrication
Panel scheme label	$panel_scheme$	Derived reporting category	0, 1, 2, 3	Index	Groups (R_p) values for interpretation, plots, tables, and dashboard filtering
Truss scheme label	$truss_scheme$	Derived reporting category	0, 1, 2	Index	Groups (R_t) values for interpretation, plots, tables, and dashboard filtering
Case-study geometry	—	Fixed input	S2 dwelling geometry	—	Defines the physical project boundary
Total number of wall panels	N_p	Fixed case-study quantity	48	Panels	Used in panel-related production and repetition calculations
Total number of roof trusses	N_t	Fixed case-study quantity	58	Trusses	Used in truss-related production and repetition calculations
Structural member sizing	—	Fixed assumption	Fixed CFS framing system	—	Not optimised in the present study
Whole-building redesign	—	Excluded	Not applicable	—	Outside the optimisation boundary

Note: panel_scheme and truss_scheme are derived reporting labels. They are not independent optimiser-controlled decision variables in the reported formulation.

The four optimiser-controlled variables should be interpreted as practical CFS framing configuration levers rather than abstract mathematical inputs. A low value of (R_p) represents a panel layout closer to the S2 reference condition, with greater panel variety and less repetition, while a high value of (R_p) represents stronger rationalisation through more repeated wall-panel families and fewer non-repeated panel-width conditions. In practical terms, increasing (R_p) would correspond to standardising wall-panel widths, reducing one-off panel families, and improving repetition in factory production and site installation. Similarly, a low value of (R_t) represents greater truss-span or truss-type diversity, while a high value represents stronger consolidation of similar roof-truss families. In practice, increasing (R_t) would involve rationalising truss layouts, group similar spans, and reducing unnecessary truss-type variation where structurally and architecturally feasible. The opening/detailing complexity variable (OC)

represents manufacturability and fit-up burden rather than the number of openings. Lower (OC) values represent more standardised lintels, jamb studs, trimmers, service penetrations, CNC cutting patterns, and site-fit-up details, whereas higher (OC) values represent more customised detailing, local reinforcement, non-repeated penetrations, and coordination burden. The prefabrication variable (P_f) represents the degree of factory preparation rather than a measured percentage of the whole building completed off site. Lower (P_f) values retain more framing work on site, while higher (P_f) values transfer more work into factory preassembly, labelling, panel completion, and site-ready framing packages. Therefore, a value such as (R_p=0.60) should not be read as a literal fabrication instruction; it represents a moderate-to-strong panel rationalisation intensity within the response model, which would later need to be translated by a designer or fabricator into actual panel-family decisions.

This interpretation is important because the optimisation searches a continuous response surface, while real CFS fabrication decisions are implemented through discrete panel widths, truss layouts, detailing standards, and factory/site work packages.

3.3.2: Continuous Rationalisation Variables and Reporting Categories

Although $R_p$and $R_t$are evaluated as continuous rationalisation variables, the model also assigns practical reporting categories to improve interpretation of the results. This is necessary because continuous rationalisation values are useful for optimisation, but construction practitioners usually interpret rationalisation in terms of scheme levels, such as reference, moderate, strong, or aggressive standardisation.

The panel scheme label is derived from the panel rationalisation intensity:

$$\text{panel\_scheme} = \text{panel\_scheme}\left(R_p\right)$$

The truss scheme label is derived from the truss rationalisation intensity:

$$\text{truss\_scheme} = \text{truss\_scheme}(R_t)$$

The continuous rationalisation variables R_p and R_t interpreted as response-model intensity variables rather than literal fabrication instructions. In practice, a fabricator would select discrete panel-width families and truss-span families. The continuous formulation is used because the optimisation framework requires a smooth and bounded design space in which rationalisation effects can be compared consistently against opening/detailing complexity and prefabrication level. The derived panel_scheme and truss_scheme categories are therefore used for reporting, interpretation, and dashboard filtering, while R_p and R_t remain the optimiser-controlled variables.

These reporting-category transformations are used only for interpretation, plotting, table generation, and dashboard filtering. They do not introduce additional degrees of freedom into the optimisation problem. The objective functions are controlled by the continuous variables R_p, R_t, OC, and P_f, while the scheme labels provide a practical way to communicate the rationalisation level selected by the optimiser.

This distinction is methodologically important. If the scheme labels were treated as additional optimiser variables, the mathematical formulation would no longer match the implemented response model. In the present formulation, the optimiser explores a continuous rationalisation response surface, and the resulting values are then grouped into construction-relevant categories.

3.3.3: Panel Rationalisation Intensity and Panel Scheme Reporting

The panel rationalisation intensity $R_p$represents the degree of wall-panel standardisation and repetition applied to the case-study CFS framing package. It is a dimensionless continuous variable bounded between 0.00 and 1.00:

$$R_p \in [0.00,1.00]$$

A value of $R_p = 0.00$represents the S2-like reference condition, while higher values represent increasing rationalisation of wall-panel production. In construction terms, increasing $R_p$represents greater standardisation of panel families, improved production repetition, and potential reductions in setup and installation effort. However, high rationalisation may also introduce penalties associated with material grouping, handling, logistics, or reduced dimensional flexibility. For this reason, the response model evaluates both benefits and penalties rather than assuming that maximum rationalisation is always optimal.

For reporting purposes, the continuous $R_p$value is grouped into a panel scheme label as follows:

$$\text{panel}_s\text{cheme}(R_p) = \begin{cases} 0, & 0.00 \le R_p < 0.20 \\ 1, & 0.20 \le R_p < 0.45 \\ 2, & 0.45 \le R_p < 0.75 \\ 3, & 0.75 \le R_p \le 1.00 \end{cases}$$

Equation 3-10

The reporting categories are summarised in Table 3-8

Table 3-8 Panel rationalisation reporting categories.

Panel scheme label	Range of (R_p)	Interpretation	Model use
0	$0.00 \le R_p < 0.20$	Reference / low panel rationalisation	S2-like or near-reference panel standardisation
1	$0.20 \le R_p < 0.45$	Mild-to-moderate panel rationalisation	Moderate grouping of panel families
2	$0.45 \le R_p < 0.75$	Strong panel rationalisation	High repetition and reduced effective panel-class diversity
3	$0.75 \le R_p \le 1.00$	Aggressive panel rationalisation	Maximum rationalisation range considered in the model

The effective number of panel classes is represented by:

$$N_{p,\text{class}}(R_p) = 4 - 2R_p$$

At the S2 reference condition:

$$N_{p,\text{class},S2} = 4$$

An intermediate value of $N_{p,\text{class}}$ does not imply that a fractional panel class exists in fabrication. It represents the production effect of progressive panel rationalisation within the continuous optimisation model.

The linear mapping is a deliberate modelling simplification. It does not imply that the commercial or fabrication effect of reducing panel classes is exactly linear in practice. A real fabricator may experience different marginal effects when moving from four panel classes to three than when moving from three to two, and similar non-linear effects may occur for truss-family rationalisation. In this study, linear mapping is used as a transparent first-order response approximation within a bounded case-study model. Potential non-linear or project-specific rationalisation effects are acknowledged as model-form uncertainty and are partly tested through the coefficient ranges, interaction terms, sensitivity analysis, and OC-fixed diagnostic results. Future work could replace this first-order mapping with fabricator-calibrated step functions or project-specific discrete scheme libraries.

3.3.4: Truss Rationalisation Intensity and Truss Scheme Reporting

The truss rationalisation intensity R_t represents the degree of roof-truss span-family rationalisation applied to the case-study CFS framing package. It is a dimensionless continuous variable bounded between 0.00 and 1.00:

$$R_t \in [0.00, 1.00]$$

A value of $R_t = 0.00$ represents the S2-like reference condition, while higher values represent increasing rationalisation of roof-truss span families. In fabrication terms, increasing R_t may reduce span-family diversity, simplify jigging, improve repetition in cutting and assembly, and support more consistent site installation. However, stronger truss rationalisation may also introduce penalties if span grouping, connection coordination, material use, or logistics become less efficient. The model therefore evaluates truss rationalisation as a trade-off variable rather than as a one-directional improvement factor.

For reporting purposes, the continuous R_t value is grouped into a truss scheme label as follows:

$$\text{truss}_\text{s}\text{cheme}(R_t) = \begin{cases} 0, & 0.00 \le R_t < 0.20 \\ 1, & 0.20 \le R_t < 0.65 \\ 2, & 0.65 \le R_t \le 1.00 \end{cases}$$

Equation 3-11

Table 3-9 Truss rationalisation reporting categories.

Truss scheme label	Range of (R_t)	Interpretation	Model use
0	$0.00 \le R_t < 0.20$	Reference / low truss rationalisation	S2-like or near-reference truss span-family diversity
1	$0.20 \le R_t < 0.65$	Moderate truss rationalisation	Partial consolidation of truss span families

Truss scheme label	Range of (R_t)	Interpretation	Model use
2	$0.65 \le R_t \le 1.00$	Aggressive truss rationalisation	High truss-family rationalisation range considered in the model

The effective number of truss types is represented by:

$$N_{t,\text{type}}(R_t) = 17 - 3R_t$$

At the S2 reference condition:

$$N_{t,\text{type},S2} = 17$$

The same effective-response interpretation used for panel classes applies to truss types.

3.4: Objective Function Vector

The optimisation problem is formulated as a three-objective minimisation problem. Each candidate CFS framing configuration is represented by the decision vector defined in the previous section:

$$\boldsymbol{x} = [R_p, R_t, OC, P_f]$$

Equation 3-12

where R_p is the panel rationalisation intensity, R_t is the truss rationalisation intensity, OC is the opening/detailing complexity index, and P_f is the prefabrication level. The variables R_p, R_t, OC, and P_f are the optimiser-controlled variables in the implemented framework. The derived panel and truss scheme labels are used later only for interpretation, plotting, and dashboard filtering.

The raw objective vector is defined as:

$$\boldsymbol{F}(\boldsymbol{x}) = [C_{\text{total}}(\boldsymbol{x}), E_{\text{CFS},A1-A5}(\boldsymbol{x}), T_{\text{process}}(\boldsymbol{x})]$$

Equation 3-13

The multi-objective optimisation problem is therefore written as:

$$\min_{\boldsymbol{x}\in\Omega} \boldsymbol{F}(\boldsymbol{x}) = \min_{\boldsymbol{x}\in\Omega}[C_{\text{total}}(\boldsymbol{x}), E_{\text{CFS},A1-A5}(\boldsymbol{x}), T_{\text{process}}(\boldsymbol{x})]$$

Equation 3-14

where Ω is the feasible design domain defined by the variable bounds for R_p, R_t, OC, and P_f. The optimisation problem is a minimisation problem because lower process cost, lower CFS framing A1–A5 carbon, and lower process-hours are preferred outcomes within the case-study boundary.

The first objective, $C_{total}(\mathbf{x})$, represents total CFS framing process cost in New Zealand dollars. The objective combines material, factory processing, site labour, setup, transport, logistics, mobilisation, waste, lifting, and overhead components within the defined CFS framing boundary.

The second objective, $E_{CFS,A1-A5}(\mathbf{x})$, represents CFS framing upfront embodied carbon within the A1–A5 system boundary. The objective includes steel production, transport, factory energy, installation, and waste-related carbon contributions as defined in the implemented response model.

The third objective, $T_{process}(\mathbf{x})$, represents total CFS framing process-hours. The objective combines factory production, setup, delivery, logistics, site installation, mobilisation, and weather-related productivity effects to provide a comparative measure of production and installation effort.

The three objective functions are formulated as comparative response models that combine fixed case-study quantities, design-variable response terms, and audited coefficients. The optimisation therefore searches the response surface defined by the implemented model, and the resulting Pareto front should be interpreted within the stated assumptions, coefficient ranges, and fixed case-study boundary.

The S2 reference case is therefore not used as a hidden optimisation constraint or embedded objective-function scaler in this section. Instead, the objective vector defines the raw model response for each candidate solution. The S2 values are applied only after raw objective evaluation to support normalisation, relative improvement calculations, and benchmark comparison.

Table 3-10 Objective units and scope.

Objective	Symbol	Unit	Included boundary	Excluded boundary	Implemented model trace
CFS framing process cost	$C_{\text{total}}(x)$	NZD	Steel material, factory processing, site labour, setup, transport, delivery/logistics, mobilisation, waste handling, lifting allowance where applicable, and overhead	Whole-house cost, land, foundations, services, linings, cladding, fit-out, professional fees, consent costs, and tender margin	Total cost response
CFS framing A1–A5 carbon	$E_{\text{CFS},A1-A5}(x)$	$kgCO_{2e}$	A1–A3 steel, A4 transport, factory energy where included, A5 installation, and waste allowance	Whole-building carbon, operational energy, maintenance, replacement, end-of-life modules, and module D recovery	A1–A5 carbon response
CFS framing process time	$T_{\text{process}}(x)$	h / process-hours	Factory production, setup, delivery/logistics, site installation, mobilisation, and weather allowance	Full CPM schedule, total calendar duration, procurement lead time, inspection delay, non-CFS trades, and subcontractor sequencing	Process-hour response

Raw objective values and normalised objective values are denoted separately throughout the study. The raw objective vector $\mathrm{F}(x)$ contains physical values in NZD, $kgCO_{2e}$, and process-hours. The normalised objective vector $\bar{\mathrm{F}}(x)$ is used only for convergence metrics, visualisation, and S2-relative comparison. This separation avoids ambiguity between physical model outputs and dimensionless optimisation diagnostics.

Table 3-11 Objective-function traceability.

Objective-function term	Model meaning	Controlled by	Fixed/reference inputs	Output
$\boldsymbol{C}_{\text{total}}(\boldsymbol{x})$	Total CFS framing process cost	R_p, R_t, OC, P_f	Baseline steel mass, labour rate, factory processing rate, setup rates, transport rate, logistics allowance, overhead allowance	NZD
$\boldsymbol{E}_{\text{CFS},\boldsymbol{A1-A5}}(\boldsymbol{x})$	CFS framing upfront carbon	R_p, R_t, OC, P_f	Steel emission factor, freight factor, transport distance, envelope area, A5 factor, waste assumption	$kgCO_{2e}$
$\boldsymbol{T}_{\textbf{process}}(\boldsymbol{x})$	CFS framing process-hours	R_p, R_t, OC, P_f	Base factory time, setup time, site installation time, delivery/logistics time, mobilisation time, weather allowance	h / process-hours

Table 3-11shows that each objective can be traced from the same decision vector to a specific model output. This traceability is important because the three objectives use different physical units and represent different parts of the CFS framing process. The framework therefore evaluates cost, carbon, and time as separate raw responses before any normalisation or ranking is applied.

The optimisation framework separates optimiser-controlled variables, derived response quantities, and fixed case-study inputs. The optimiser directly controls R_p, R_t, OC, and P_f. Derived quantities include effective panel and truss counts, reporting categories, repetition measures, material-response terms, and objective-function components. Fixed inputs include the case-study geometry, wall-panel count, roof-truss count, steel section, steel grade, coating class, and S2 reference quantities. This separation provides traceability between the decision vector, intermediate response terms, and final cost, carbon, and process-hour outputs

3.5: Cost Objective

The cost objective evaluates CFS framing process cost for each candidate configuration x = $[R_p, R_t, OC, P_f]$ in this R_p the panel rationalisation intensity, R_t is the truss rationalisation intensity, OC is the opening/detailing complexity index, and P_f is the prefabrication level. The cost objective is governed by these four design variables through derived response terms such as material weight, factory processing demand, setup effort, site installation effort, logistics burden, opening/detailing response, and waste allowance.

The total CFS framing process cost is defined as:

$$C_{\text{total}}(x) = C_{\text{direct}}(x) + C_{\text{oh}}(x)$$

Equation 3-15

where $C_{\text{total}}(\mathbf{x})$ is the total process cost in NZD, $C_{\text{direct}}(\mathbf{x})$ is the direct CFS framing cost, and $C_{\text{oh}}(\mathbf{x})$ is the overhead and coordination allowance applied within the CFS framing scope.

The direct cost subtotal is calculated as:

$$\begin{aligned} C_{\text{direct}}(x) = {} & C_{\text{steel}}(x) + C_{\text{fac}}(x) + C_{\text{site}}(x) + C_{\text{setup}}(x) \\ & + C_{\text{trans}}(x) + C_{\text{log}}(x) + C_{\text{mob}}(x) + C_{\text{open}}(x) \\ & + C_{\text{waste}}(x) + C_{\text{lift}}(x) \end{aligned}$$

Equation 3-16

Therefore, the complete cost objective may also be written as:

$$\begin{aligned} C_{\text{total}}(x) = {} & C_{\text{steel}}(x) + C_{\text{fac}}(x) + C_{\text{site}}(x) + C_{\text{setup}}(x) \\ & + C_{\text{trans}}(x) + C_{\text{log}}(x) + C_{\text{mob}}(x) + C_{\text{open}}(x) \\ & + C_{\text{waste}}(x) + C_{\text{lift}}(x) + C_{\text{oh}}(x) \end{aligned}$$

Equation 3-17

The material cost is calculated from the predicted CFS steel mass:

$$C_{\text{steel}}(x) = W(x)\, c_{\text{steel}}$$

Equation 3-18

where $W(\mathbf{x})$is the modelled CFS steel mass in kg and $c_{\text{steel}}$is the steel material rate in NZD/kg. The steel mass is a derived response term rather than an independent optimiser variable.

The factory processing cost is calculated as:

$$C_{\text{fac}}(x) = W(x)\, c_{\text{fac}}\, M_{\text{fac}}(x)$$

Equation 3-19

where $c_{\text{fac}}$is the factory processing rate in NZD/kg and $M_{\text{fac}}(\mathbf{x})$is the factory-processing response multiplier. This multiplier represents the effect of rationalisation and prefabrication on factory-side effort.

The site labour cost is calculated from the site installation time:

$$C_{\text{site}}(x) = T_{\text{site}}(x)\, c_{\text{lab}}$$

Equation 3-20

where $T_{\text{site}}(\mathbf{x})$is the site installation time in hours and $c_{\text{lab}}$is the labour rate in NZD/hr. This term links the cost objective to the process-time response model because changes in rationalisation, opening/detailing complexity, and prefabrication level can alter the amount of site installation effort required.

Setup, jigging, and CNC preparation cost are represented as:

$$C_{\text{setup}}(x) = C_{\text{setup,base}}\, M_{\text{setup}}(x)$$

Equation 3-21

where $C_{\text{setup,base}}$is the S2 reference setup allowance and $M_{\text{setup}}(\mathbf{x})$is the setup response multiplier. The setup multiplier captures the effect of rationalisation on production preparation while retaining a lower-bound logic so that setup cost cannot be reduced unrealistically to zero.

The reference setup allowance is calculated as:

$$C_{\text{setup,base}} = N_{p,\text{class},S2}\, c_{p,\text{setup}} + N_{t,\text{type},S2}\, c_{t,\text{setup}} + C_{\text{jig}}$$

Equation 3-22

where $N_{p,\text{class},S2}$is the S2 effective panel-class count, $N_{t,\text{type},S2}$is the S2 effective truss-type count, $c_{p,\text{setup}}$is the setup cost per panel class, $c_{t,\text{setup}}$is the setup cost per truss type, and $C_{\text{jig}}$is the base jigging and CNC preparation allowance.

The transport cost is calculated from transported CFS mass, distance, and freight rate:

$$C_{\text{trans}}(x) = \left(\frac{W(x)}{1000}\right) d_{\text{trans}} c_{\text{trans}} M_{\text{trans}}(x)$$

Equation 3-23

The term $W(x)/1000$ converts CFS framing mass from kilograms to tonnes. The unit structure is therefore $(\text{kg}/1000) \times \text{km} \times \text{NZD}/\text{t-km} = \text{NZD}$, before applying the dimensionless transport modifier $M_{\text{trans}}(x)$..

Delivery and logistics cost are represented as:

$$C_{\text{log}}(x) = T_{\text{log}}(x)\, c_{\text{log}}$$

Equation 3-24

where $T_{\text{log}}(\mathbf{x})$is the delivery and logistics time allowance in hours and $c_{\text{log}}$is the logistics cost rate in NZD/hr. This term accounts for truck scheduling, unloading coordination, sequencing, and site-handling effort.

Site mobilisation cost is calculated as:

$$C_{\text{mob}}(x) = T_{\text{mob}}\, c_{\text{mob}}$$

Equation 3-25

where $T_{\text{mob}}$is the mobilisation time allowance and $c_{\text{mob}}$is the mobilisation cost rate. In the present case-study model, mobilisation is treated as a fixed CFS framing setup activity rather than a variable building-wide programme item.

Opening/detailing cost is represented as the combined cost effect of opening standardisation and opening complexity:

$$C_{\text{open}}(x) = C_{\text{open,std}}(x) + C_{\text{open,comp}}(x)$$

Equation 3-26

where $C_{\text{open,std}}(\mathbf{x})$represents the cost response associated with standardised opening/detailing conditions and $C_{\text{open,comp}}(\mathbf{x})$represents the cost penalty associated with increased detailing complexity. This term captures lintels, jamb studs, trimming, local reinforcement, CNC cutting, service-penetration coordination, and site fit-up effects. It does not add or remove windows or doors.

Waste handling cost is calculated as:

$$C_{\text{waste}}(x) = \left(\frac{W(x) r_{\text{waste,ref}}}{1000}\right) c_{\text{waste}}$$

Equation 3-27

Where $r_{\text{waste,ref}}$ is the steel waste rate and c_{waste} is the waste handling rate in NZD/t. The waste rate is treated as a response term because rationalisation, prefabrication, and opening/detailing complexity can affect offcuts, rework, and material-handling burden.

The lifting allowance is written as:

$$C_{\text{lift}}(x) = C_{\text{lift,allow}}$$

Equation 3-28

where $C_{\text{lift,allow}}$is the lifting equipment allowance. In the current case-study calculation this value is set to zero because no separate crane or telehandler cost evidence is applied.

The overhead component is applied to the direct CFS framing cost subtotal:

$$C_{\text{oh}}(x) = r_{\text{oh}}\, C_{\text{direct}}(x)$$

Equation 3-29

where $r_{\text{oh}}$is the overhead fraction. This term represents coordination and indirect cost allowance within the CFS framing scope. It is not a whole-project commercial margin.

Table 3-12 Cost-rate and response-coefficient audit for the CFS framing cost objective.

Cost item	Symbol	Central value	Unit	Function in cost objective	Source / basis
Steel material rate	c_{steel}	2.40	NZD/kg	Converts CFS steel mass into material cost	NZ market benchmark / case-study assumption
Factory processing rate	c_{fac}	2.80	NZD/kg	Converts CFS steel mass and factory multiplier into processing cost	CFS processing allowance / case-study assumption
Site labour rate	c_{lab}	48.00	NZD/hr	Converts site installation time into labour cost	Labour rate plus on-cost allowance
Panel setup rate	$c_{p,\text{setup}}$	25.00	NZD/class	Represents panel setup, jigging, and CNC preparation burden	Setup allowance

Cost item	Symbol	Central value	Unit	Function in cost objective	Source / basis
Truss setup rate	$c_{t,\text{setup}}$	35.00	NZD/type	Represents truss setup, jigging, and CNC preparation burden	Setup allowance
Base jig/CNC allowance	C_{jig}	800.00	NZD	Represents baseline production preparation allowance	Case-study setup assumption
Transport rate	c_{trans}	3.00	NZD/t-km	Converts transported steel mass and distance into freight cost	NZ freight benchmark assumption
Transport distance	d_{trans}	121.00	km	Defines the assumed supply-chain distance	Auckland–Waikato scenario assumption
Delivery/logistics rate	$c_{\log}$	48.00	NZD/hr	Converts logistics time into coordination cost	Set equal to site labour rate
Delivery/logistics base time	$T_{\log,S2}$	2.00	h	Reference truck scheduling and unloading coordination allowance	S2 process assumption
Site mobilisation rate	c_{mob}	48.00	NZD/hr	Converts mobilisation time into cost	Set equal to site labour rate
Site mobilisation time	T_{mob}	3.00	h	Represents CFS framing site setup and mobilisation	S2 process assumption
Reference waste rate	r_{waste}	0.05	fraction	Estimates steel offcut and waste handling quantity	Case-study waste allowance
Waste handling rate	c_{waste}	250.00	NZD/t	Converts waste steel mass into handling cost	Waste handling/disposal allowance
Lifting allowance	$C_{\text{lift,allow}}$	0.00	NZD	Optional lifting-equipment allowance	No separate lifting evidence applied
Overhead fraction	r_{oh}	0.10	fraction	Applies coordination/overhead allowance to direct CFS framing cost	Framing-scope overhead assumption
Logistics cost response coefficient	$k_{\log,C}$	0.110	dimensionless	Adjusts logistics burden under rationalisation and prefabrication response	Calibrated response coefficient
Low opening/detailing cost coefficient	$k_{\text{OC,low}}$	0.030	dimensionless	Represents cost response associated with standardised opening/detailing conditions	Calibrated response coefficient
High opening/detailing cost coefficient	$k_{\text{OC,high}}$	0.220	dimensionless	Represents cost penalty for increased opening/detailing complexity	Calibrated response coefficient
Factory prefabrication cost coefficient	$k_{\text{prefab,fac}}$	0.20	dimensionless	Represents increased factory-side cost response with higher prefabrication	Calibrated response coefficient
Prefabrication overload cost coefficient	$k_{\text{prefab,over}}$	0.42	dimensionless	Represents nonlinear cost penalty at high prefabrication levels	Calibrated response coefficient
Prefabrication transport-response coefficient	$k_{P_f,\text{trans}}$	0.06	dimensionless	Transport response assumption	Adjusts transport cost with prefabrication level

Cost item	Symbol	Central value	Unit	Function in cost objective	Source / basis
Panel transport-response coefficient	$k_{R_p,\mathrm{trans}}$	0.025	dimensionless	Transport response assumption	Transport response assumption
Truss transport-response coefficient	$k_{R_t,\mathrm{trans}}$	0.02	dimensionless	Transport response assumption	Adjust transport cost with R_t
Setup lower-bound factor	$m_{\mathrm{setup,min}}$	0.45	fraction	Physical setup cannot reduce to zero	Prevents unrealistic collapse of setup cost
Panel setup penalty coefficient	$k_{R_p,\mathrm{setup,pen}}$	0.4	dimensionless	Excessive panel rationalisation penalty	Adds setup burden above rationalisation knee
Truss setup penalty coefficient	$k_{R_t,\mathrm{setup,pen}}$	0.3	dimensionless	Excessive truss rationalisation penalty	Adds setup burden above rationalisation knee

The numerical coefficients used in cost, carbon, time, and interaction response equations are reported through coefficient-audit tables rather than treated as hidden model inputs. Each coefficient is documented with its value, unit, source or basis, confidence level, and model use. The source/basis field distinguishes between project-derived quantities, model-calculated reference values, literature-supported rates, published database values, scenario assumptions, and calibrated response coefficients. The confidence level is not a statistical confidence interval; it is an engineering transparency rating that indicates the strength of evidence supporting each parameter. High-confidence parameters are directly project-derived or based on well-established published factors. Medium-confidence parameters are supported by literature, industry benchmarks, or transparent scenario logic but remain project-dependent. Low-confidence or medium-low-confidence parameters are retained as bounded assumptions where direct CFS project evidence was not available. This classification allows the examiner and future users to distinguish reliable fixed boundary data from assumptions that may require recalibration for another project.

In engineering terms, the cost objective is structured so that steel mass controls material cost, factory processing captures production effort, site labour captures installation effort, setup terms capture the burden of unique panel and truss families, and logistics terms capture transport and handling. The strongest assumptions are the calibrated setup, logistics, opening/detailing, and prefabrication-response coefficients, because these are project-dependent and not directly measured from commercial factory records. If measured supplier pricing or factory productivity data became available, these coefficients could be recalibrated while retaining the same objective-function structure.

3.6: Carbon Objective

The second objective function evaluates the CFS framing upfront carbon associated with each candidate design vector X in that R_p is the panel rationalisation intensity, R_t is the truss rationalisation intensity, OC is the opening/detailing complexity index, and P_f is the prefabrication level. The carbon objective is governed by these variables through derived response terms such as steel mass, transport burden, factory production time, waste factor, opening/detailing response, and prefabrication-related A5 saving.

Increasing prefabrication level affects the carbon objective through opposing mechanisms. A higher P_f increases the factory-side contribution because more framing work is transferred to factory processing. At the same time, it reduces the site-side A5 installation contribution by lowering site handling, rework, weather exposure, and installation activity. Within the coefficient ranges used in this model, the A5 reduction is larger than the additional factory-energy contribution, so the net response of P_f is carbon-reducing across the tested design range. This is interpreted as a model-structure finding within the fixed CFS framing A1–A5 boundary, not as universal empirical proof that higher prefabrication always reduces carbon.

The CFS framing carbon objective is limited to the A1–A5 upfront-carbon boundary. It is written as:

$$E_{\text{CFS},A1-A5}(x) = E_{A1-A3}(x) + E_{A4}(x) + E_{\text{fac}}(x) + E_{A5}(x) + E_{\text{waste}}(x)$$

Equation 3-30

where $E_{\text{CFS},A1-A5}(\mathbf{x})$is the total CFS framing A1–A5 carbon in $\text{kgCO}_{2\text{e}}$. The term E_{A1-A3} represents product-stage steel carbon, E_{A4} represents transport carbon, E_{fac} represents factory energy carbon where included, E_{A5} represents construction-stage installation carbon, and E_{waste} represents steel offcut or waste allowance where included.

The A1–A3 steel carbon is calculated from the modelled CFS steel mass:

$$E_{A1-A3}(x) = W(x)\, EF_{\text{steel}}$$

Equation 3-31

where $W(\mathbf{x})$is the modelled CFS steel mass in kg and $EF_{\text{steel}}$is the A1–A3 steel emission factor in $\text{kgCO}_{2\text{e}}$ /kg. In the current model, the central steel emission factor is $2.10\ \text{kgCO}_2\text{e/kg}$. This value is treated as a benchmark factor for early-stage decision support rather than as a product-specific environmental product declaration.

The A4 transport carbon is calculated using transported steel mass, transport distance, freight emission factor, and a transport-response multiplier:

$$E_{A4}(x) = \left(\frac{W(x)}{1000}\right) d_{\text{trans}} EF_{\text{freight}} M_{A4}(x)$$

Equation 3-32

The unit structure is $(\mathrm{kg}/1000) \times \mathrm{km} \times \mathrm{kgCO_2e/t\text{-}km} = \mathrm{kgCO_2e}$. For the S2 reference transport scenario, $d_{\text{trans}} = 121$ km and $EF_{\text{freight}} = 0.135$ $\mathrm{kgCO_{2e}}$/t-km, giving an equivalent factor of 0.016335 $\mathrm{kgCO_{2e}}$ /kg.

The A4 transport factor is treated as a scenario-based carbon coefficient rather than a measured field value. For the reported S2 reference case, the assumed transport distance is 121 km, and the freight carbon factor is 0.135 $\mathrm{kgCO_{2e}}$/t-km. This gives an equivalent transport carbon factor of $EF_{A4} = 0.016335$ $\mathrm{kgCO_{2e}}$/kg of transported CFS framing steel. The A4 factor is therefore location-sensitive and recalibrated for projects with different fabrication locations, transport distances, freight modes, or component weights.

$$EF_{A4} = \frac{121 \times 0.135}{1000} = 0.016335\ \mathrm{kgCO_2e/kg}$$

Equation 3-33

The transport-response multiplier is represented as:

$$M_{A4}(x) = 1 + k_{\log,E}L(x) + k_{R_p,A4}R_p + k_{R_t,A4}R_t$$

Equation 3-34

where $L(\mathbf{x})$is the logistics response term, $k_{\log,E}$is the logistics-carbon coefficient, $k_{R_p,A4}$is the panel rationalisation transport-carbon coefficient, and $k_{R_t,A4}$is the truss rationalisation transport-carbon coefficient. This term allows higher prefabrication and rationalisation to influence transport-related carbon where handling and freight burden increase.

Factory energy carbon is included when factory energy is activated in the model:

$$E_{\text{fac}}(x) = T_{\text{fac,prod}}(x)P_{\text{fac}}EF_{\text{grid}}$$

Equation 3-35

where $T_{\text{fac,prod}}(\mathbf{x})$is the factory production time in hours, $P_{\text{fac}}$is the assumed average factory equipment power in kW, and $EF_{\text{grid}}$is the electricity emission factor in $\mathrm{kgCO_{2e}}$/kWh. In the current model, factory energy is included using an average factory equipment power of 5.0 kW and an electricity emission factor of $0.08\ \mathrm{kgCO_2e/kWh}$. Factory energy is therefore treated as an explicit model term rather than being hidden inside a single global steel factor.

The A5 installation carbon is calculated using the envelope/framing area and an A5 installation factor:

$$E_{A5}(x) = A_e EF_{A5} M_{A5}(x)$$

Equation 3-36

where $A_e$is the envelope/framing area in m², $EF_{A5}$is the A5 installation emission factor in $\mathrm{kgCO_{2e}}$/m², and $M_{A5}(\mathbf{x})$is the A5 response multiplier. The current model uses $A_e = 217.2555\ \mathrm{m}^2$and $EF_{A5} = 3.20\ \mathrm{kgCO_2e/m^2}$

The A5 response multiplier is represented as:

$$M_{A5}(x) = 1 - k_{P_f,A5}B_{P_f}(x) + k_{OC,E}OC_{\text{high}}(x) + k_{OC,\text{low},E}OC_{\text{low}}(x)$$

Equation 3-37

where $B_{P_f}(\mathbf{x})$is the prefabrication benefit response, $k_{P_f,A5}$is the prefabrication-related A5 carbon saving coefficient, $OC_{\text{high}}(\mathbf{x})$is the high opening/detailing complexity response, and $OC_{\text{low}}(\mathbf{x})$is the low-complexity detailing response. This formulation reflects the model assumption that higher prefabrication can reduce site-stage carbon, while complex openings and detailing can increase installation-stage burden.

To avoid unrealistic collapse of construction-stage carbon, the A5 component is bounded below:

$$E_{A5}(x) \geq 0.35A_e EF_{A5}$$

Equation 3-38

This lower bound recognises that even a highly prefabricated framing option still requires unloading, handling, positioning, fixing, connection work, site coordination, and installation support. Therefore, A5 carbon cannot be reduced to zero.

Waste carbon is calculated from the modelled steel mass, steel emission factor, and a variable waste response factor:

$$E_{\text{waste}}(x) = W(x)r_{\text{waste}}(x)EF_{\text{steel}}$$

Equation 3-39

where $r_{\text{waste}}(\mathbf{x})$is the waste response factor. In contrast to the cost objective, where waste handling uses the reference waste rate, the carbon objective allows the waste factor to vary with rationalisation, prefabrication, and opening/detailing complexity. In the current model, the waste factor is bounded to avoid unrealistic values.

The waste response factor is represented as:

$$r_{\text{waste}}(x) = r_{\text{waste,ref}}\left[1 - k_{R_p,\text{waste}}R_p - k_{R_t,\text{waste}}R_t - k_{P_f,\text{waste}}B_{P_f}(x) + k_{OC,\text{waste}}OC_{\text{high}}(x)\right]$$

Equation 3-40

where $r_{\text{waste,ref}}$is the reference steel waste rate, $k_{R_p,\text{waste}}$and $k_{R_t,\text{waste}}$represent rationalisation-related waste savings, $k_{P_f,\text{waste}}$represents prefabrication-related waste reduction, and $k_{OC,\text{waste}}$represents the increase in waste burden under complex opening/detailing conditions. In the current model, $r_{\text{waste,ref}} = 0.05$. The carbon waste factor is clipped between $0.015$and 0.080, reflecting a bounded steel offcut allowance rather than an unrestricted response.

The complete carbon objective is therefore:

$$
\begin{aligned}
E_{\mathrm{CFS},A1-A5}(x) = {} & W(x)EF_{\mathrm{steel}} \\
& + \left(\frac{W(x)}{1000}\right) d_{\mathrm{trans}} EF_{\mathrm{freight}} M_{A4}(x) \\
& + T_{\mathrm{fac,prod}}(x) P_{\mathrm{fac}} EF_{\mathrm{grid}} \\
& + A_e EF_{A5} M_{A5}(x) \\
& + W(x) r_{\mathrm{waste}}(x) EF_{\mathrm{steel}}
\end{aligned}
$$

Equation 3-41

Separate site energy is not added as an additional term in the current model. This avoids double counting because the A5 installation allowance already represents construction-stage site activity, equipment, fuel, waste, and installation burden. The model therefore includes factory energy explicitly where activated but treats site-stage energy through the A5 installation term.

Table 3-13 Carbon coefficient audit.

Carbon item	**symbol**	**Central value**	**Unit**	**Source / evidence basis**	**Uncertainty treatment / model use**
A1–A3 steel emission factor	EF_{steel}	2.1	$kgCO_{2e}$/kg	BRANZ / ICE / EPD-type steel benchmark	Tested using steel carbon factor range
Freight emission factor	EF_{freight}	0.135	$kgCO_{2e}$/t-km	NZ road-freight emission factor assumption	Tested using freight emission factor range
Transport distance	d_{trans}	121	km	Auckland–Waikato delivery-distance assumption	Can be varied as scenario parameter
Equivalent A4 factor	$EF_{A4,\mathrm{kg}}$	0.016335	$kgCO_{2e}$/kg	Scenario-based A4 transport factor derived from 121 km and 0.135 $kgCO_{2e}$/t-km	Used to calculate freight-related CFS framing A4 carbon
A5 installation factor	EF_{A5}	3.2	$kgCO_{2e}$/m²	Residential A5 installation-carbon allowance	Tested using A5 carbon factor range
Envelope/framing area	A_e	217.2555	m²	Case-study geometry	Used in A5 installation carbon
Electricity emission factor	EF_{grid}	0.08	$kgCO_{2e}$/kWh	NZ electricity grid factor assumption	Used for factory energy carbon
Factory average equipment power	P_{fac}	5	kW	Factory equipment-power assumption range (3-10)	Used for factory energy carbon
Factory energy inclusion switch	—	TRUE	Boolean	Factory energy included in current model	Included as explicit factory-energy term
Separate site energy inclusion switch	—	FALSE	Boolean	Avoids double counting with A5 installation allowance	Separate site energy excluded
Waste carbon inclusion switch	—	TRUE	Boolean	Waste carbon included in current model	Includes steel offcut/waste carbon
Reference steel waste rate	$r_{\mathrm{waste,ref}}$	0.05	fraction	Case-study offcut allowance	Used as base waste factor
Panel waste-saving coefficient	$k_{R_p,\mathrm{waste}}$	0.012	dimensionless	Rationalisation waste response	Reduces waste factor as R_p increases
Truss waste-saving coefficient	$k_{R_t,\mathrm{waste}}$	0.007	dimensionless	Rationalisation waste response	Reduces waste factor as R_t increases

Carbon item	symbol	Central value	Unit	Source / evidence basis	Uncertainty treatment / model use
Prefabrication waste-reduction coefficient	$k_{P_f,\text{waste}}$	0.15	dimensionless	Prefabrication waste response	Reduces waste factor with prefabrication benefit
Logistics carbon coefficient	$k_{\log,E}$	0.04	dimensionless	Logistics carbon response	Adjusts A4 transport carbon
Panel transport-carbon coefficient	$k_{R_p,A4}$	0.025	dimensionless	Panel rationalisation transport response	Adjusts A4 transport carbon
Truss transport-carbon coefficient	$k_{R_t,A4}$	0.02	dimensionless	Truss rationalisation transport response	Adjusts A4 transport carbon
Prefabrication A5 saving coefficient	$k_{P_f,A5}$	0.35	dimensionless	Prefabrication site-stage carbon response	Reduces A5 installation carbon with prefabrication
High opening/detailing carbon coefficient	$k_{OC,E}$	0.12	dimensionless	Opening/detailing complexity response	Increases A5 carbon under complex detailing
Low opening/detailing A5 coefficient	$k_{OC,\text{low},E}$	0.015	dimensionless	Standardised opening/detailing response	Small A5 response for low-complexity detailing
Waste high-complexity coefficient	$k_{OC,\text{waste}}$	0.04	dimensionless	Opening/detailing waste response	Increases waste factor under high opening complexity
Waste-factor lower bound	$r_{\text{waste,min}}$	0.015	fraction	Prevents unrealistically low waste carbon	Applied to waste response factor
Waste-factor upper bound	$r_{\text{waste,max}}$	0.08	fraction	Prevents unrealistically high waste carbon	Applied to waste response factor
A5 lower-bound factor	$m_{A5,min}$	0.35	fraction	Site installation carbon cannot reduce to zero	Applied to A5 installation carbon
Factory average power	P_{fac}	5.0	kW	3–10 range between this	Scenario-based allowance for small CFS roll-forming/assembly energy use; bounded because factory energy is a minor carbon term

Table 3-13 separates direct emission factors, scenario inputs, fixed case-study quantities, boundary switches, and calibrated response coefficients used in the carbon objective.

In engineering terms, the carbon objective is dominated by steel mass and the A1–A3 steel emission factor, while A4 transport, factory energy, A5 installation, and waste allowances provide smaller but traceable contributions. The strongest carbon assumption is the use of a benchmark steel emission factor rather than a product-specific EPD for the exact CFS product. If project-specific EPDs, measured transport records, or verified A5 construction-stage data became available, the same model structure could be updated by replacing the relevant coefficients.

3.7: Time Objective

The third objective function evaluates the CFS framing process time for each candidate in same design vector x where R_p is the panel rationalisation intensity, R_t is the truss rationalisation intensity, OC is the opening/detailing complexity index, and P_f is the prefabrication level. These variables influence the time objective through factory production effort, setup requirements, delivery/logistics, site installation time, opening/detailing effects, rationalisation penalties, prefabrication benefits, and weather exposure.

The total CFS framing process time is defined as:

$$T_{\text{process}}(x) = T_{\text{fac,prod}}(x) + T_{\text{setup}}(x) + T_{\text{log}}(x) + T_{\text{mob}}(x) + T_{\text{site}}(x) + T_{\text{weather}}(x)$$

Equation 3-42

where $T_{\text{process}}(\mathbf{x})$ is measured in process-hours. The term $T_{\text{fac,prod}}$ represents factory production time, T_{setup} represents factory setup and CNC/jig preparation, T_{log} represents delivery and logistics time, T_{mob} represents site mobilisation, T_{site} represents site installation time, and T_{weather} represents the site weather-exposure allowance.

Factory production time is represented as:

$$T_{\text{fac,prod}}(x) = T_{\text{fac,prod},S2} M_{\text{fac,time}}(x)$$

Equation 3-43

where $T_{\text{fac,prod},S2}$ is the S2 reference factory production time and $M_{\text{fac,time}}(\mathbf{x})$ Is the factory-time response multiplier. In the current model, factory production time is affected by prefabrication level, panel and truss rationalisation, opening/detailing complexity, and nonlinear factory overload at high prefabrication levels. Higher prefabrication can increase factory-side effort, while rationalisation can reduce repeated production effort up to a point. Excessive rationalisation or complex opening/detailing may add fabrication and coordination burden.

The S2 factory panel and truss production times are treated as fixed case-study reference coefficients rather than measured universal productivity rates. Because project-specific factory production records were not available, $T_{\text{fac,panel,S2}} = 11.70\text{h}$ and $T_{\text{fac,truss,S2}} = 24.20\text{h}$ is reported with ±20% plausibility ranges in the coefficient audit. These ranges communicate coefficient confidence and engineering uncertainty, but the base times are not independently sampled as Monte Carlo variables in the reported robustness analysis.

Factory setup time is calculated as:

$$T_{\text{setup}}(x) = T_{\text{setup},S2} M_{\text{setup,time}}(x)$$

Equation 3-44

where $T_{\text{setup},S2}$is the S2 reference setup time and $M_{\text{setup,time}}(\mathbf{x})$is the setup-time response multiplier. Setup time decreases as panel and truss rationalisation increase because fewer repeated families reduce preparation, jigging, and CNC-programming effort. However, this reduction is bounded because setup work cannot be eliminated completely.

The setup-time lower-bound condition is:

$$T_{\text{setup}}(x) \geq 0.40T_{\text{setup},S2}$$

Equation 3-45

This lower bound prevents unrealistic collapse of setup time. Even highly rationalised framing still requires model checking, jigging, machine preparation, labelling, handling, and quality-control activities.

Site installation time is represented as:

$$T_{\text{site}}(x) = T_{\text{site},S2}M_{\text{site,time}}(x)$$

Equation 3-46

where $T_{\text{site},S2}$ is the S2 reference site installation time and $M_{\text{site,time}}(\mathbf{x})$ is the site-time response multiplier. Site installation time is reduced by prefabrication benefit and rationalisation effects, but it is increased by logistics burden, excessive rationalisation penalties, and high opening/detailing complexity. Low opening/detailing complexity may reduce site fit-up time because lintels, jambs, trimmers, service penetrations, and connection details are more standardised.

A physical lower bound is applied to site installation time:

$$T_{\text{site}}(x) \geq T_{\text{site,min}}$$

Equation 3-47

where $T_{\text{site,min}}$ is the minimum feasible site handling and installation time. In the current case-study model, this value is 28.0 h. This recognises that a CFS framing package still requires unloading, positioning, fixing, alignment, connections, and inspection-related handling even when prefabrication is high.

$$T_{\text{log}}(x) = T_{\text{log},S2}M_{\text{log,time}}(x)$$

Equation 3-48

where $T_{\text{log},S2}$ is the S2 reference delivery/logistics time and $M_{\text{log,time}}(\mathbf{x})$is the logistics-time response multiplier. This term captures truck scheduling, unloading windows, sequencing, handling coordination, and additional logistics burden caused by higher prefabrication and larger or more rationalised framing elements.

Site mobilisation time is treated as a fixed CFS framing allowance:

$$T_{\text{mob}}(x) = T_{\text{mob},S2}$$

Equation 3-49

where $T_{\text{mob},S2}$represents the reference site setup and mobilisation time. In the current case-study model, site mobilisation is not treated as a variable project-wide scheduling activity. It is retained as a fixed CFS framing process allowance.

Weather exposure is represented as a site-time allowance:

$$T_{\text{weather}}(x) = r_{\text{weather}} s_{\text{weather}} T_{\text{site}}(x)$$

Equation 3-50

where $r_{\text{weather}}$is the weather downtime fraction and $s_{\text{weather}}$is the weather shielding factor. The current model uses a Waikato weather downtime fraction of 0.10and a shielding factor of 0.60. Weather exposure is applied to site installation time only because factory production is assumed to occur under controlled conditions.

$$T_{\text{weather}} = D_{\text{weather}}\, S_{\text{weather}}\, T_{\text{site}}$$

Equation 3-51

$$D_{\text{weather}} = 0.10, \qquad S_{\text{weather}} = 0.60$$

The complete time objective can therefore be written as:

$$\begin{aligned} T_{\text{process}}(x) = {} & T_{\text{fac,prod},S2} M_{\text{fac,time}}(x) + T_{\text{setup},S2} M_{\text{setup,time}}(x) \\ & + T_{\text{log},S2} M_{\text{log,time}}(x) + T_{\text{mob},S2} + T_{\text{site},S2} M_{\text{site,time}}(x) \\ & + r_{\text{weather}} s_{\text{weather}} T_{\text{site}}(x) \end{aligned}$$

Equation 3-52

This formulation evaluates the total CFS framing effort in process-hours. It does not model full project scheduling, activity dependencies, crew calendars, inspection waiting time, procurement lead time, or overlapping work by other trades.

Table 3-14 Time coefficient audit.

Item	Symbol	Central value	Unit	Source / evidence basis	Uncertainty treatment / model use
Factory truss production time	$T_{\text{fac,truss},S2}$	24.20	h	S2 process assumption Range (19.36–29.04) h	Forms part of reference factory production time
Factory panel production time	$T_{\text{fac,panel},S2}$	11.70	h	S2 process assumption Range (9.36–14.04) h	Forms part of reference factory production time
Factory production time	$T_{\text{fac,prod},S2}$	35.90	h	Sum of truss and panel factory production time	Used as base for factory production response
Factory setup time	$T_{\text{setup},S2}$	3.50	h	S2 jig/CNC/setup process assumption	Modified by setup-time response multiplier
Delivery/logistics base time	$T_{\text{log},S2}$	2.00	h	S2 truck scheduling and unloading allowance	Modified by logistics-time response multiplier
Site mobilisation time	$T_{\text{mob},S2}$	3.00	h	S2 site setup and mobilisation assumption	Treated as fixed CFS framing allowance
Truss installation time	$T_{\text{site,truss},S2}$	18.18	h	S2 installation process assumption	Forms part of reference site installation time
Panel installation time	$T_{\text{site,panel},S2}$	57.06	h	S2 installation process assumption	Forms part of reference site installation time
Site installation time	$T_{\text{site},S2}$	75.24	h	Sum of panel and truss installation time	Used as base for site installation response
Minimum site installation time	$T_{\text{site,min}}$	28.00	h	Physical lower-bound assumption for handling 48 panels and 58 trusses	Prevents unrealistic collapse of site time
Setup lower-bound factor	$m_{\text{setup,min}}$	0.40	fraction	Physical setup cannot reduce to zero	Lower bound applied to setup time

Item	Symbol	Central value	Unit	Source / evidence basis	Uncertainty treatment / model use
Weather downtime fraction	r_{weather}	0.10	fraction	Waikato weather exposure scenario	Tested using weather downtime range
Weather shielding factor	s_{weather}	0.60	fraction	Prefabrication/weather-exposure scenario assumption	Tested using shielding-factor range
Factory prefabrication time coefficient	$k_{P_f,\text{fac,time}}$	0.28	dimensionless	Prefabrication shifts work into factory production	Increases factory production time with P_f
Site prefabrication saving coefficient	$k_{P_f,\text{site,time}}$	0.60	dimensionless	Prefabrication reduces site installation effort	Reduces site time through prefabrication benefit
Panel site-saving coefficient	$k_{R_p,\text{site}}$	0.030	dimensionless	Panel rationalisation/repetition response	Reduces site installation time with R_p
Truss site-saving coefficient	$k_{R_t,\text{site}}$	0.020	dimensionless	Truss rationalisation/repetition response	Reduces site installation time with R_t
Panel factory penalty coefficient	$k_{R_p,\text{fac}}$	0.050	dimensionless	Excess rationalisation can increase factory burden	Used in factory-time response
Truss factory penalty coefficient	$k_{R_t,\text{fac}}$	0.030	dimensionless	Excess rationalisation can increase factory burden	Used in factory-time response
Low opening/detailing factory-time coefficient	$k_{OC,\text{low,fac}}$	0.030	dimensionless	Standardised detailing still requires factory coordination	Adds small factory-side time response
High opening/detailing time coefficient	$k_{OC,\text{high,time}}$	0.280	dimensionless	Complex openings increase fabrication and site fit-up burden	Increases site-time response
Factory overload time coefficient	$k_{P_f,\text{over,time}}$	0.30	dimensionless	High prefabrication can create factory bottleneck effects	Adds nonlinear time penalty above saturation threshold
Prefabrication saturation threshold	$P_{f,\text{sat}}$	0.78	fraction	Factory-capacity response assumption	Activates nonlinear overload response
Logistics time coefficient	$k_{\log,T}$	0.150	dimensionless	Higher prefabrication and rationalisation can increase coordination burden	Adjusts delivery/logistics and site-time response
Panel logistics-time coefficient	$k_{R_p,\text{log,T}}$	0.030	dimensionless	Panel rationalisation can affect handling and sequencing	Used in delivery/logistics response
Truss logistics-time coefficient	$k_{R_t,\text{log,T}}$	0.020	dimensionless	Truss rationalisation can affect handling and sequencing	Used in delivery/logistics response
Panel setup saving coefficient	$k_{R_p,\text{setup,T}}$	0.22	dimensionless	Panel rationalisation reduces setup repetition burden	Reduces factory setup time
Truss setup saving coefficient	$k_{R_t,\text{setup,T}}$	0.12	dimensionless	Truss rationalisation reduces setup repetition burden	Reduces factory setup time
Panel setup penalty coefficient	$k_{R_p,\text{setup,pen}}$	0.45	dimensionless	Excessive panel rationalisation can add grouping/jig burden	Adds setup-time penalty above rationalisation knee

Item	Symbol	Central value	Unit	Source / evidence basis	Uncertainty treatment / model use
Truss setup penalty coefficient	$k_{R_t,\text{setup,pen}}$	0.35	dimensionless	Excessive truss rationalisation can add grouping/jig burden	Adds setup-time penalty above rationalisation knee
Panel site penalty coefficient	$k_{R_p,\text{site,pen}}$	0.18	dimensionless	Excessive panel rationalisation can increase site handling burden	Adds site-time penalty above rationalisation knee
Truss site penalty coefficient	$k_{R_t,\text{site,pen}}$	0.14	dimensionless	Excessive truss rationalisation can increase site handling burden	Adds site-time penalty above rationalisation knee
Low opening/detailing site-time saving coefficient	$k_{OC,\text{low,site}}$	0.08	dimensionless	Standardised detailing can reduce site fit-up effort	Reduces site installation time under low (OC)
High opening/detailing factory-time coefficient	$k_{OC,\text{high,fac}}$	0.120	dimensionless	High detailing complexity increases factory coordination	Adds factory-time penalty
weather shielding factor	(S_{weather})	0.6	fraction	Scenario factor representing partial shielding/reduction of site exposure through prefabrication	Reduces the weather-exposed portion of site installation time
weather downtime fraction	(D_{weather})	0.1	fraction	Scenario allowance for Waikato site-exposed CFS installation activity	Weather allowance in process-hours objective

In engineering terms, the time objective treats CFS framing as a factory–site process rather than a single installation-duration estimate. Factory production, setup, delivery/logistics, mobilisation, site installation, and weather allowance are modelled separately so that prefabrication can increase factory effort while reducing site effort. The most uncertain assumptions are productivity, weather, logistics, and prefabrication-saving coefficients. If measured factory and site time records became available, these coefficients could be recalibrated without changing the decision vector.

3.8: Weight and Interaction Model

The interaction model introduces response terms that represent how rationalisation, opening/detailing complexity, prefabrication level, logistics, waste, and steel mass interact within the case-study boundary. These interaction terms are not claimed to be universal empirical laws. Several coefficients are engineering-informed assumptions because project-specific CFS fabrication data at this level of detail were not publicly available. For this reason, Table 3-15 identifies the basis and confidence of each interaction coefficient, while Chapter 6 tests the influence of these assumptions through sensitivity analysis, Monte Carlo uncertainty propagation, and diagnostic checks. The optimisation therefore searches for an explicitly disclosed response model rather than a hidden or externally validated production model.

The objective functions defined in Sections 3.5 to 3.7 depend on intermediate response terms rather than on the four design variables directly. The design vector is: $\boldsymbol{x} = [R_p, R_t, OC, P_f]$

where R_p is the panel rationalisation intensity, R_t is the truss rationalisation intensity, OC is the opening/detailing complexity index, and P_f is the prefabrication level. These variables affect the objectives through derived model responses including steel mass, effective repetition, setup effort, factory burden, site productivity, transport/logistics burden, waste, and A5 installation response.

The effective panel-class and truss-type quantities defined in Section 3.3 are used here only to calculate repetition response terms. They do not replace the underlying case-study panel and truss schedules and do not represent structural redesign.

The steel-weight response is defined as:

$$W(\boldsymbol{x}) = W_{S2} f_W(R_p, R_t, OC, P_f)$$

Equation 3-53

where $W(\mathbf{x})$is the modelled CFS steel mass in kg, $W_{S2}$is the S2 reference steel mass, and $f_W$is the steel-weight response multiplier. The reference mass is:

$$W_{S2} = 2195.264 \text{ kg}$$

The weight multiplier is written as:

$$\begin{aligned} f_W = 1 &- k_{R_p,w}R_p - k_{R_t,w}R_t - k_{P_f,w}\Delta P_n^+ \\ &+ k_{R_p,\text{over}}G_p + k_{R_t,\text{over}}G_t \\ &+ k_{OC,\text{low},w}OC_{\text{low}} + k_{OC,\text{high},w}OC^*_{\text{high}} \end{aligned}$$

Equation 3-54

where $k_{R_p,w}$and $k_{R_t,w}$represent repetition-related waste or offcut savings from panel and truss rationalisation. The term $k_{P_f,w}\Delta P_n^+$represents a small prefabrication precision saving above the S2 prefabrication reference. The terms $G_p$and $G_t$represent nonlinear oversizing or grouping penalties caused by excessive rationalisation. The terms $OC_{\text{low}}$and OC^*_{high}represent the steel effect of low and high opening/detailing complexity.

The normalised prefabrication level is:

$$P_n = \frac{P_f - P_{f,min}}{P_{f,max} - P_{f,min}}$$

Equation 3-55

The positive prefabrication deviation above the S2 reference is:

$$\Delta P_n^+ = \max(0, P_n - P_{n,S2})$$

Equation 3-56

The rationalisation oversizing terms are defined as:

$$G_p = \max\left(0, R_p - R_{p,\mathrm{knee}}\right)^2$$

$$G_t = \max\left(0, R_t - R_{t,\mathrm{knee}}\right)^2$$

Equation 3-57

The rationalisation knee point $R_{p,\mathrm{knee}} = R_{t,\mathrm{knee}} = 0.65$ represents the transition from beneficial standardisation to over-rationalisation. Below this threshold, reducing panel or truss family variation is treated mainly as a repetition and setup benefit. Beyond this threshold, further rationalisation may require oversizing, trimming, additional handling, or less efficient grouping of components. The nonlinear penalty therefore prevents the model from assuming that maximum rationalisation is always beneficial.

Non-repeated openings can affect adjacent studs, lintels, trimmers, CNC cutting, labelling, and site alignment simultaneously. For this reason, high opening complexity is represented through OC_{high} and a compounded penalty term rather than through a simple one-to-one percentage increase.

The range $OC \in [0.80, 1.30]$ represents a practical DfMA/detailing spectrum from highly standardised repeatable opening modules to heavily customised non-repeated detailing. The upper bound does not mean that every cost, carbon, or time component increases by 30%. Instead, $OC$is a bounded response-model input whose effects are applied through specific opening/detailing coefficients affecting trimming, lintels, jamb studs, reinforcement, CNC cutting, service penetrations, and site fit-up.

The opening/detailing response terms are:

$$OC_{\mathrm{low}} = \max(0, 1 - OC)$$

Equation 3-58

$$OC_{\mathrm{high}} = \max(0, OC - 1)$$

Equation 3-59

$$OC^{*}_{\mathrm{high}} = OC_{\mathrm{high}} + OC^{2}_{\mathrm{high}}$$

Equation 3-60

Low opening/detailing complexity represents greater standardisation of lintels, jamb studs, trimmers, service penetrations, and CNC cutting patterns. This can reduce site fit-up effort, but it may still require standardised detailing modules and therefore is not treated as a free removal of material. High opening/detailing complexity represents custom openings, non-repeated trimmers, local reinforcement, coordination burden, and additional site fit-up.

Because the model uses a fixed structural section and fixed case-study geometry, steel mass is bounded as:

$$0.94W_{S2} \leq W(x) \leq 1.06W_{S2}$$

Equation 3-61

3.8.1: Engineering Basis for Response and Interaction Coefficients

The coefficient set is a central part of the methodology because it defines how the four decision variables influence steel mass, factory effort, logistics, site labour, A5 construction-stage effects, and penalty or saving mechanisms. The coefficients were not introduced as arbitrary tuning factors. They were selected to represent bounded engineering mechanisms within the fixed CFS case-study boundary, including repetition benefit, rationalisation penalty, prefabrication saturation, opening/detailing burden, logistics effects, and mass-related carbon response.

Table 3-15 therefore acts as the coefficient audit for the implemented response model. It brings together the interaction coefficients used to connect (R_p), (R_t), (OC), and (P_f) to material, cost, carbon, time, factory, logistics, and site-response terms. The table should be read as a methodological traceability table rather than as a set of universal industry constants. The values define the calibrated response structure used in this case-study model and are later tested through sensitivity and Monte Carlo robustness analysis.

Table 3-15 Engineering justification and uncertainty treatment for response and interaction coefficients.

Item	Symbol	Central value	Unit	Source / evidence basis	Uncertainty treatment / model use
S2 steel mass	W_{S2}	2195.264	kg	Direct case-study material quantity	Fixed reference mass for weight response
Weight lower bound	$m_{W,min}$	0.94	fraction	Fixed section and fixed geometry assumption	Prevents unrealistic mass reduction
Weight upper bound	$m_{W,max}$	1.06	fraction	Fixed section and fixed geometry assumption	Prevents unrealistic mass increase
Panel waste/mass saving coefficient	$k_{R_p,w}$	0.012	dimensionless	Panel repetition/offcut saving assumption	Reduces mass and waste response with R_p
Truss waste/mass saving coefficient	$k_{R_t,w}$	0.007	dimensionless	Truss repetition/offcut saving assumption	Reduces mass and waste response with R_t
Prefabrication precision saving coefficient	$k_{P_f,w}$	0.012	dimensionless	Improved factory cutting and handling precision	Applies only above S2 prefabrication reference
Panel oversizing coefficient	$k_{R_p,\text{over}}$	0.160	dimensionless	Excessive panel rationalisation/grouping assumption	Adds mass, cost, and handling penalty above knee
Truss oversizing coefficient	$k_{R_t,\text{over}}$	0.120	dimensionless	Excessive truss rationalisation/grouping assumption	Adds mass, cost, and handling penalty above knee
Panel rationalisation knee	$R_{p,\text{knee}}$	0.65	dimensionless	Rationalisation response threshold	Activates nonlinear oversizing/grouping penalty
Truss rationalisation knee	$R_{t,\text{knee}}$	0.65	dimensionless	Rationalisation response threshold	Activates nonlinear oversizing/grouping penalty
Low opening/detailing	$k_{OC,\text{low},w}$	0.020	dimensionless	Standardised lintel/jamb/detail module assumption	Adds small material response under low (OC)

Item	Symbol	Central value	Unit	Source / evidence basis	Uncertainty treatment / model use
material coefficient					
High opening/detailing material coefficient	$k_{OC,\mathrm{high},w}$	0.060	dimensionless	Extra trimming/local reinforcement assumption	Adds material response under high (OC)
Panel setup saving coefficient	$k_{R_p,\mathrm{setup}}$	0.22	dimensionless	Fewer panel families reduce setup burden	Reduces setup cost and setup time
Truss setup saving coefficient	$k_{R_t,\mathrm{setup}}$	0.12	dimensionless	Fewer truss families reduce setup burden	Reduces setup cost and setup time
Setup lower-bound factor	$m_{\mathrm{setup,min}}$	0.45 cost / 0.40 time	fraction	Setup cannot be eliminated physically	Prevents setup collapse in cost/time response
Panel setup penalty coefficient	$k_{R_p,\mathrm{setup,pen}}$	0.40 cost / 0.45 time	dimensionless	Excessive panel grouping/jigging penalty	Adds setup burden above rationalisation knee
Truss setup penalty coefficient	$k_{R_t,\mathrm{setup,pen}}$	0.30 cost / 0.35 time	dimensionless	Excessive truss grouping/jigging penalty	Adds setup burden above rationalisation knee
Panel factory penalty coefficient	$k_{R_p,\mathrm{fac}}$	0.050	dimensionless	Factory burden under rationalised panel families	Used in factory cost and factory time
Truss factory penalty coefficient	$k_{R_t,\mathrm{fac}}$	0.030	dimensionless	Factory burden under rationalised truss families	Used in factory cost and factory time
Logistics cost coefficient	$k_{\mathrm{log},C}$	0.110	dimensionless	Handling, delivery, unloading, and sequencing burden	Tested through Monte Carlo coefficient range
Logistics time coefficient	$k_{\mathrm{log},T}$	0.150	dimensionless	Coordination and site-handling burden	Tested through Monte Carlo coefficient range
Logistics carbon coefficient	$k_{\mathrm{log},E}$	0.040	dimensionless	Transport and handling carbon response	Tested through Monte Carlo coefficient range
Factory prefabrication time coefficient	$k_{P_f,\mathrm{fac,time}}$	0.28	dimensionless	Prefabrication shifts work into factory	Increases factory production time
Factory prefabrication cost coefficient	$k_{P_f,\mathrm{fac}}$	0.20	dimensionless	Higher prefabrication increases factory processing burden	Increases factory processing cost
Site prefabrication saving coefficient	$k_{P_f,\mathrm{site}}$	0.60	dimensionless	Off-site work reduces site installation effort	Tested through prefabrication time-saving range
A5 prefabrication saving coefficient	$k_{P_f,A5}$	0.35	dimensionless	Higher prefabrication reduces site-stage carbon burden	Tested through A5 carbon-saving range
Prefabrication saturation threshold	$P_{f,\mathrm{sat}}$	0.78	fraction	Factory capacity / overload threshold	Activates nonlinear overload response
Prefabrication overload cost coefficient	$k_{P_f,\mathrm{over},C}$	0.42	dimensionless	High-prefabrication factory overload assumption	Adds nonlinear factory cost penalty
Prefabrication overload time coefficient	$k_{P_f,\mathrm{over},T}$	0.30	dimensionless	High-prefabrication factory bottleneck assumption	Adds nonlinear factory time penalty

Item	Symbol	Central value	Unit	Source / evidence basis	Uncertainty treatment / model use
Prefabrication overload carbon coefficient	$k_{P_f,\text{over},E}$	0.08	dimensionless	High-prefabrication factory burden assumption	Retained as response coefficient for carbon interpretation
High opening/detailing cost coefficient	$k_{OC,\text{high},C}$	0.220	dimensionless	Complex openings increase trimming and coordination	Used in factory cost and opening cost response
High opening/detailing time coefficient	$k_{OC,\text{high},T}$	0.280	dimensionless	Complex openings increase site and factory time	Used in process-time response
High opening/detailing carbon coefficient	$k_{OC,\text{high},E}$	0.120	dimensionless	Complex openings increase A5/material burden	Used in carbon response
Low opening/detailing factory-time coefficient	$k_{OC,\text{low,fac}}$	0.030	dimensionless	Standardisation still requires factory coordination	Adds small factory-time response
Low opening/detailing site-time saving coefficient	$k_{OC,\text{low,site}}$	0.080	dimensionless	Standardised detailing improves site fit-up	Reduces site installation time

The steel-mass response coefficients were selected to keep material variation within a narrow and physically defensible range around the S2 reference mass. This is necessary because the optimiser does not redesign the CFS section, member capacity, connections, bracing, or compliance checks. The lower and upper mass bounds therefore prevent unrealistic material reductions or increases while still allowing rationalisation, prefabrication precision, waste, and detailing effects to influence the cost and carbon objectives.

The rationalisation coefficients for (R_p) and (R_t) were selected to represent the competing effects of repetition and grouping. Increased panel or truss rationalisation can reduce setup burden, improve repetition, reduce waste, and simplify production. However, excessive rationalisation can also introduce oversizing, grouping penalties, local inefficiency, and coordination burden. The model therefore includes both saving and penalty mechanisms so that rationalisation is not automatically rewarded at the upper boundary.

The prefabrication coefficients were selected to represent the shift of work from site installation to factory-controlled preparation. Higher (P_f) can reduce site labour, site fit-up, and A5 site-stage burden, but it can also increase factory processing, handling, logistics coordination, and overload effects at high prefabrication levels. The implemented response therefore uses prefabrication benefit and penalty terms to represent diminishing returns rather than assuming that maximum prefabrication is always optimal.

The opening/detailing complexity coefficients were selected to represent the additional fabrication and site burden associated with openings, lintels, jambs, trimmers, local reinforcement, angle plates, and fit-up requirements. Lower (OC) represents simplified or more standardised detailing relative to S2, while higher (OC) represents increased detailing burden.

This structure allows the model to capture the practical effect of opening/detailing decisions without treating (OC) as a direct window-area ratio or an automatically extracted CAD metric.

The logistics, factory, and site interaction coefficients were selected to connect the decision variables to the process mechanisms affected by prefabricated CFS framing. Rationalisation and prefabrication can reduce some setup and site processes, but they can also increase transport, handling, coordination, and factory-stage effort. These coefficients allow the model to represent the factory–site trade-off that is central to the research problem.

The coefficient values should be interpreted as calibrated response-model assumptions rather than measured universal constants. Their role is to provide a transparent, auditable, and bounded representation of the case-study response structure. The sensitivity and Monte Carlo analyses in Chapter 6 test the effect of selected coefficient uncertainty, but they do not constitute validation of all possible response forms. External validation would require measured cost, carbon, fabrication, logistics, and site-productivity data from multiple CFS

The interaction model also affects setup, factory cost, site time, logistics, and carbon. Setup decreases with rationalisation but is penalised when rationalisation becomes excessive:

$$M_{\text{setup}}(x) = \max\left[m_{\text{setup,min}}, 1 - k_{R_p,\text{setup}}R_p - k_{R_t,\text{setup}}R_t + k_{R_p,\text{setup,pen}}G_p + k_{R_t,\text{setup,pen}}G_t\right]$$

Equation 3-62

Separate lower-bound values are used for the cost and time setup response terms. The cost setup multiplier uses $m_{\text{setup,min},C} = 0.45$, while the time setup multiplier uses $m_{\text{setup,min},T} = 0.40$. This distinction reflects that setup cost and setup time do not necessarily reduce at the same rate under rationalisation. This term reflects the engineering logic that fewer repeated panel and truss families reduce setup effort, but over-rationalisation may introduce additional jigging, grouping, checking, and coordination burden.

The logistics response is represented as:

$$L(x) = P_n^2\left(0.65 + k_{R_p,\text{log}}R_p + k_{R_t,\text{log}}R_t\right)$$

Equation 3-63

This term captures the practical effect that high prefabrication can increase transport, sequencing, unloading, and site-handling coordination. Panel and truss rationalisation may further increase this effect where larger or more integrated elements are produced.

The factory processing response is represented as:

$$\begin{aligned} M_{\text{fac}}(x) = 1 &+ k_{P_f,\text{fac}}P_n + k_{R_p,\text{fac}}R_p + k_{R_t,\text{fac}}R_t \\ &+ k_{P_f,\text{over}}O_f + k_{OC,\text{low}}OC_{\text{low}}^2 + k_{OC,\text{high}}OC_{\text{high}}^* \end{aligned}$$

Equation 3-64

where O_f is the nonlinear factory overload response activated above the prefabrication saturation threshold:

$$O_f = \left[\frac{\max(0, P_f - P_{f,\text{sat}})}{P_{f,max} - P_{f,\text{sat}}}\right]^2$$

Equation 3-65

This reflects the assumption that high prefabrication can reduce site effort, but may also increase factory-side processing, scheduling, handling, and production-load effects.

The site-time response uses the same interaction logic. Prefabrication and rationalisation can reduce site effort, while excessive rationalisation, logistics, and high opening/detailing complexity can increase it:

$$\begin{aligned} M_{\text{site}}(x) = 1 &- k_{P_f,\text{site}} B_{P_f} - k_{R_p,\text{site}} R_p - k_{R_t,\text{site}} R_t \\ &+ k_{R_p,\text{site,pen}} G_p + k_{R_t,\text{site,pen}} G_t \\ &- k_{OC,\text{low,site}} OC_{\text{low}} + k_{OC,\text{high,time}} OC^*_{\text{high}} + k_{\log,T} L(x) \end{aligned}$$

Equation 3-66

The prefabrication benefit term $B_{P_f}$is a saturating response rather than a linear assumption:

$$B_{P_f} = \frac{1 - \exp(-2.3P_n)}{1 - \exp(-2.3)}$$

Equation 3-67

The response-model structure itself is also a modelling assumption. In particular, the rationalisation knee at ($R_{p,knee} = R_{t,knee} = 0.65$) and the prefabrication benefit exponent of 2.3 define the location and curvature of the benefit–penalty transition in the implemented model. These terms were selected to represent diminishing returns, over-rationalisation effects, and the practical limit beyond which additional prefabrication or standardisation no longer produces proportional site-effort savings. However, they were not calibrated against measured multi-project CFS fabrication or site-installation records. Therefore, the reported Pareto front should be interpreted as conditional on both the selected coefficient values and the selected response-model form. The sensitivity and Monte Carlo analyses in Chapter 6 test uncertainty in coefficient values and selected variable responses, but they do not constitute full validation of alternative response structures. Future work should therefore test alternative rationalisation knee locations, prefabrication saturation exponents, and interaction forms using measured factory productivity, site installation, logistics, and project-specific cost/carbon data.

3.9: Normalisation and S2 Improvement Metrics

The cost, carbon, and process-time objectives are first calculated as raw physical quantities: $C_{\text{total}}(\mathbf{x})$in NZD, $E_{\text{CFS},A1-A5}(\mathbf{x})$in $\text{kgCO}_{2\text{e}}$, and $T_{\text{process}}(\mathbf{x})$in process-hours. Because these objectives have different units and magnitudes, they are normalised against the S2 reference values before comparison, plotting, ranking, and dashboard interpretation.

The S2 reference values are:

$$C_{S2} = 18335.47 \text{ NZD}$$

$$E_{S2} = 5366.79 \text{ kgCO}_2\text{e}$$

$$T_{S2} = 93.92 \text{ h}$$

For any candidate solution $\mathbf{x}$, the normalised objectives are defined as:

$$\bar{C}(\boldsymbol{x}) = \frac{C_{\text{total}}(\boldsymbol{x})}{C_{S2}}$$

$$\bar{E}(\boldsymbol{x}) = \frac{E_{\text{CFS},A1-A5}(\boldsymbol{x})}{E_{S2}}$$

$$\bar{T}(\boldsymbol{x}) = \frac{T_{\text{process}}(\boldsymbol{x})}{T_{S2}}$$

From Equation 3-5

where $\bar{C}(\mathbf{x})$, $\bar{E}(\mathbf{x})$, and $\bar{T}(\mathbf{x})$are dimensionless normalised values. At the S2 reference condition:

$$\bar{C}(\boldsymbol{x}_{\boldsymbol{S2}}) = \bar{E}(\boldsymbol{x}_{\boldsymbol{S2}}) = \bar{T}(\boldsymbol{x}_{\boldsymbol{S2}}) = 1$$

A value below 1.00 indicates that the candidate solution is lower than S2 for that objective. A value above 1.00 indicates that the candidate solution is higher than S2.

For clearer reporting, the normalised values are also converted into signed percentage-change metrics. The cost percentage change relative to S2 is:

$$\Delta C_{\%}(\boldsymbol{x}) = \left(\frac{C_{\text{total}}(\boldsymbol{x})}{C_{S2}} - 1\right) \times 100$$

Equation 3-68

The carbon percentage change relative to S2 is:

$$\Delta E_{\%}(\boldsymbol{x}) = \left(\frac{E_{\text{CFS},A1-A5}(\boldsymbol{x})}{E_{S2}} - 1\right) \times 100$$

Equation 3-69

The process-time percentage change relative to S2 is:

$$\Delta T_{\%}(\boldsymbol{x}) = \left(\frac{T_{\text{process}}(\boldsymbol{x})}{T_{S2}} - 1\right) \times 100$$

Equation 3-70

The sign convention is important. A negative value indicates improvement because the objective is lower than S2. A positive value indicates a penalty because the objective is higher than S2. Therefore, $\Delta C_{\%} = -5$means a 5% cost reduction relative to S2, while $\Delta C_{\%} =$

$+5$means a 5% cost increase relative to S2. The same interpretation applies to carbon and process time.

The signed percentage-change convention is retained throughout the results chapter, representative solution tables, Monte Carlo robustness outputs, and dashboard views. This avoids changing the meaning of the metric between sections. Where the wording "saving" or "reduction" is used, it refers only to negative percentage-change values.

Table 3-16 Normalisation and sign convention.

Item	Symbol	Equation / value	Unit	Interpretation	Model use
S2 cost denominator	C_{S2}	18335.47	NZD	Model-calculated S2 CFS framing cost	Cost normalisation denominator
S2 carbon denominator	E_{S2}	5366.79	$kgCO_{2e}$	Model-calculated S2 CFS framing A1–A5 carbon	Carbon normalisation denominator
S2 time denominator	T_{S2}	93.92	h	Model-calculated S2 CFS framing process time	Time normalisation denominator
Normalised cost	$\bar{C}(x)$	$C_{\text{total}}(x)/C_{S2}$	dimensionless	Below 1.00 means lower cost than S2	Pareto plotting, ranking, dashboard
Normalised carbon	$\bar{E}(x)$	$E_{\text{CFS},A1-A5}(x)$ $/E_{S2}$	dimensionless	Below 1.00 means lower carbon than S2	Pareto plotting, ranking, dashboard
Normalised time	$\bar{T}(x)$	$T_{\text{process}}(x)$ $/T_{S2}$	dimensionless	Below 1.00 means lower process time than S2	Pareto plotting, ranking, dashboard
Cost percentage change	$\Delta C_{\%}$	$(C/C_{S2}-1)$ $\times 100$	%	Negative = cost reduction; positive = cost increase	S2 comparison table
Carbon percentage change	$\Delta E_{\%}$	$(E/E_{S2}-1)$ $\times 100$	%	Negative = carbon reduction; positive = carbon increase	S2 comparison table
Time percentage change	$\Delta T_{\%}$	$(T/T_{S2}-1)$ $\times 100$	%	Negative = time reduction; positive = time increase	S2 comparison table
S2 reference condition	$\bar{C}=\bar{E}$ $=\bar{T}=1$	$at(x_{S2})$	dimensionless	S2 is the comparison baseline	Benchmark reference

The multi-objective results are interpreted using the normalisation and improvement metric proposed in this section. Raw objective values are kept in physical units, normalised values enable comparisons across cost, carbon and time. When using signed percentage changes, interpretation will be relative to S2, so that negative numbers will signify a decrease relative to S2, and positive numbers a rise relative to S2. This convention is adhered throughout the results of the optimisation, typical solution tables, the uncertainty analysis, and the dashboard interpretation.

3.10: Pareto Dominance and S2-Relative Improvement

The multi-objective optimisation has three objectives: CFS framing process cost, CFS framing A1-A5 carbon, and CFS framing process time. In the above-mentioned goal, these objectives are to be minimised simultaneously and therefore the optimisation does not look for a single optimisation solution with one universal best score. Rather, it finds nondominated alternatives corresponding to various cost–carbon–time trade-offs. This follows the Pareto-ranking principle to be used in NSGA-II which sorts the (candidate) solutions by nondominance prior to selection and reproduction (Deb et al., 2002).

For two candidate solutions, $\mathbf{x}_a$and $\mathbf{x}_b$, let the objective vector be:

$$F(x) = \left[C_{\text{total}}(x), E_{\text{CFS},A1-A5}(x), T_{\text{process}}(x)\right]$$

Equation 3-71

Solution $\mathbf{x}_a$is said to dominate solution $\mathbf{x}_b$if they are no worse in all objectives and strictly better in at least one objective:

$$x_a \prec x_b \Leftrightarrow f_i(x_a) \leq f_i(x_b)\ \forall i \in \{C, E, T\} \quad \text{and} \quad \exists j: f_j(x_a) < f_j(x_b)$$

Equation 3-72

where f_i represents one of the three objective functions. A candidate solution is nondominated if no other feasible candidate solution dominates it. The set of all nondominated solutions forms the Pareto front:

$$\mathcal{P} = \{x \in \Omega: \nexists y \in \Omega \text{ such that } y \prec x\}$$

Equation 3-73

where $\mathcal{P}$ is the Pareto set and $\Omega$is the feasible design domain. A solution on the Pareto front cannot be improved in one objective without worsening at least one other objective within the evaluated design space.

Pareto dominance between candidate solutions is different from dominance relative to S2. S2 is not an optimiser-selected solution; it is the model-calculated benchmark used for comparison. Therefore, a Pareto solution is reported as dominating S2 only when all three raw objectives are lower than the corresponding S2 benchmark values:

$$C_{\text{total}}(x) < C_{S2}$$

$$E_{\text{CFS},A1-A5}(x) < E_{S2}$$

$$T_{\text{process}}(x) < T_{S2}$$

The strict S2 Improvement and Reference-Case Interpretation condition is therefore:

$$x \prec x_{S2} \Leftrightarrow C_{\text{total}}(x) < C_{S2} \wedge E_{\text{CFS},A1-A5}(x) < E_{S2} \wedge T_{\text{process}}(x) < T_{S2}$$

Equation 3-74

This is a strict definition for the three objectives to be minimization objectives and hence suitable for research interpretation. Solutions that only lower cost; only lower carbon; only lower time, or only two out of the three are not seen as dominating S2. It is rather reported as a trade-off option. This difference is significant, because the existing implementation differentiates between representative solutions in terms of improvement of all three goals and between solutions in which only some goals are improved and solutions which are primarily analytical (diagnostic). Also included in the response model are s2-dominance checks based on a cost, carbon and time comparison with the s2 benchmark values.

Pareto dominance between candidate solutions is different from S2-relative improvement. S2 is not treated as an optimiser-selected Pareto solution; it is the reference case used for comparison. Therefore, a candidate solution is reported as an all-objective improvement relative to S2 only when all three raw objective values are lower than the corresponding S2 benchmark values.

The S2 benchmark values used in this comparison are:

$C_{S2} = 18335.47$ NZD, $E_{S2} = 5366.79$ kgCO_2e, $T_{S2} = 93.92$ h

These values are used only after the raw objective functions have been evaluated. They do not modify the raw cost, carbon, or time equations.

Table 3-17 Classification of Pareto solutions relative to the S2 benchmark.

Result condition	Mathematical condition	Classification	Research reporting wording
Better than S2 in all three objectives	$(C < C_{S2})$, $(E < E_{S2})$, and $(T < T_{S2})$	S2-relative improvement	"The solution improves upon S2 across all three modelled objectives. because it reduces cost, carbon, and process time simultaneously."
Better in cost and carbon only	$(C < C_{S2})$, $(E < E_{S2})$, $(T \geq T_{S2})$	Trade-off solution	"The solution reduces cost and carbon but requires a process-time penalty."
Better in cost and time only	$(C < C_{S2})$, $(T < T_{S2})$, $(E \geq E_{S2})$	Trade-off solution	"The solution reduces cost and process time but increases carbon."
Better in carbon and time only	$(E < E_{S2})$, $(T < T_{S2})$, $(C \geq C_{S2})$	Trade-off solution	"The solution reduces carbon and process time but requires a cost premium."
Better in one objective only	Only one of (C), (E), or (T) is below S2	Specialist solution	"The solution is objective-specific and should not be presented as a broad improvement over S2."
Not better in any objective	$(C \geq C_{S2})$, $(E \geq E_{S2})$, and $(T \geq T_{S2})$	Dominated or diagnostic comparison point	"The solution does not improve on S2 and is retained only for diagnostic or comparison purposes."
Pareto-optimal but not S2-dominant	Solution is nondominated within Ω, but at least one objective is worse than S2	Pareto trade-off solution	"The solution is Pareto-optimal within the search space but does not dominate S2."

This table is important for Chapter 5 because it prevents overclaiming. A Pareto-front solution is only described as dominating S2 when all three raw objectives are lower than their S2 values; otherwise, it is interpreted as a trade-off solution relative to the benchmark.

The dominance definitions in this section establish the interpretation rules used later in the results chapter. Pareto dominance identifies nondominated trade-off solutions within the feasible design space, whereas S2 Improvement and Reference-Case Interpretation evaluates whether those solutions improve on the model-calculated benchmark. This distinction prevents overstatement of the optimisation results. A solution may be Pareto-optimal because it represents an efficient trade-off, but it is reported as dominating S2 only when cost, A1–A5 carbon, and process time are all lower than the S2 reference values. Solutions that improve only one or two objectives are therefore interpreted as trade-off alternatives rather than complete benchmark improvements.

3.11: Constraints and Feasibility Screening

The optimisation framework applies feasibility screening to ensure that candidate CFS framing configurations remain within a practical engineering boundary. The model uses a fixed structural system and fixed CFS section, while the optimiser varies only the configuration-level variables R_p, R_t, OC, and P_f. The purpose of the constraint system is to maintain feasible and practically constructible CFS framing configurations within the defined case-study boundary.

The fixed CFS section used in the case-study model is: $89S41/0.75/G500/Z275$

where the steel thickness is 0.75mm, the steel grade is G500, and the coating class is Z275. This section is treated as a fixed project boundary condition. The optimiser does not vary section depth, thickness, steel grade, coating, member spacing, or structural capacity.

The feasible design domain is written as: $x \in \Omega$

where $\Omega$is the feasible CFS framing design domain after applying hard constraints and soft constructability screening. The decision vector remains: $x = \left[R_p, R_t, OC, P_f\right]$

The design-variable bounds are:

$$0 \le R_p \le 1$$

$$0 \le R_t \le 1$$

$$0.80 \le OC \le 1.30$$

$$0.50 \le P_f \le 0.90$$

Equation 3-75

These bounds define the configuration space explored by the optimisation. They do not change the building geometry, CFS section, or structural design basis. In the implemented framework, these limits are applied directly to the decoded individual before objective evaluation, so candidate solutions remain within the intended design-variable range.

The hard feasibility constraints are applied to practical CFS framing limits, including maximum panel width, maximum panel height, stud spacing, steel thickness, and estimated panel lift mass. A candidate solution that violates a hard constraint is treated as infeasible through a large penalty in the objective space. The hard feasibility conditions are:

$$b_{\text{panel}} \leq 2.40 \text{ m}$$

$$h_{\text{panel}} \leq 3.00 \text{ m}$$

$$s_{\text{stud}} \leq 600 \text{ mm}$$

$$t_{\text{steel}} \geq 0.55 \text{ mm}$$

$$m_{\text{lift}} \leq 2000 \text{ kg}$$

where $b_{\text{panel}}$is the maximum panel width, $h_{\text{panel}}$is the assumed panel height, $s_{\text{stud}}$is the stud spacing, $t_{\text{steel}}$is the steel thickness, and $m_{\text{lift}}$is the estimated panel lift mass. These constraints represent constructability screening checks applied within the optimisation framework.

This distinction is important because AS/NZS 4600:2018 sets the minimum requirements for the design of cold-formed steel structural members. The present optimisation framework does not perform AS/NZS 4600 sizing or capacity checks. Instead, it assumes that the fixed CFS framing system has already been selected and uses feasibility screening to keep the optimisation within a practical case-study boundary (Standards Australia / Standards New Zealand, 2018).

Soft constructability constraints are also applied. These do not automatically invalidate a solution, but they add penalties where the design becomes less desirable from a production or site-execution perspective. The soft screening conditions include a maximum preferred number of unique truss spans, a minimum factory repetition score, and a maximum preferred opening/detailing complexity:

$$N_{t,\text{unique}} \leq 25$$

$$R_{\text{factory}} \geq 0.65$$

$$OC \leq 1.20$$

where $N_{t,\text{unique}}$is the number of unique truss span families, $R_{\text{factory}}$is the factory repetition score, and $OC$is the opening/detailing complexity index. These soft constraints reflect constructability preference rather than absolute physical impossibility.

The penalty-adjusted fitness values are represented as:

$$F_C^*(\boldsymbol{x}) = \bar{C}(\boldsymbol{x}) + P_C(\boldsymbol{x})$$

$$F_E^*(\boldsymbol{x}) = \bar{E}(\boldsymbol{x}) + P_E(\boldsymbol{x})$$

$$F_T^*(\boldsymbol{x}) = \bar{T}(\boldsymbol{x}) + P_T(\boldsymbol{x})$$

where F_C^*, F_E^*, and $F_T^*$are the constrained fitness values used during optimisation, and P_C, P_E, and $P_T$are the cost, carbon, and time penalties associated with hard or soft constraint violations. Hard constraint violations receive a very large penalty, while soft constraint violations receive smaller proportional penalties. The implemented framework follows this structure by adding separate costs, carbon, and time penalties to the normalised objective values.

Additional fixed case-study screening conditions are applied to ensure that the evaluated configurations remain within the intended CFS framing boundary:

$$w_{p,max} \leq 2.4$$

$$h_{p,max} \leq 3.0$$

$$s_{stud} \leq 600 \text{ mm}$$

$$t_{CFS} = 0.75 \text{ mm} \geq 0.55 \text{ mm}$$

$$M_{lift,max} \leq 2000\text{kg}$$

Only feasible solutions are retained for Pareto-front interpretation and dashboard export. Constraint screening is therefore used to preserve the practical case-study boundary rather than to perform detailed structural code verification.

The constraint-handling approach keeps the optimisation within a practical CFS framing decision space. Hard constraints prevent physically unrealistic panel, section, spacing, and lifting conditions, while soft constraints represent constructability preferences related to truss repetition, factory repetition, opening/detailing complexity, and optional carbon-intensity scenarios. The constraint-handling strategy maintains optimisation within a practical CFS framing decision space by combining hard constructability limits with soft preference-based screening conditions.

3.12: NSGA-II with LHS Initialisation

Section 3.12 defines the implemented LHS-initialised NSGA-II workflow. The general rationale for NSGA-II and LHS was established in Sections 2.7 and 2.8. This section therefore reports only the thesis-specific configuration: decision variables, bounds, initialisation, objective evaluation, feasibility screening, selection settings, convergence diagnostics, grid-reference comparison, multi-seed stability checking, and nondominated archive extraction.

Each candidate solution is represented by a continuous design vector $x = [R_p, R_t, OC, P_f]$ in which the optimisation variables are constrained within predefined feasible bounds, as specified in Equation (3-75).

The normalised objective functions are evaluated from the corresponding raw cost, carbon, and time formulations presented in Sections 3.5 to 3.9. During fitness evaluation, constraint penalties are incorporated into the normalised objective values so that infeasible or undesirable configurations are systematically penalised during the search process.

The bounded crossover and bounded mutation are used to generate new candidate solutions. Crossover mixes the continuous variables of parent solutions, and mutation helps to add controlled perturbations within the bounds. Once they are reproduced, this population including the parents is ranked again with nondominated sorting and crowding distance. To recover population size, candidates are ranked and the best and different slate of candidates are chosen to survive into the next generation.

The number of generations is predetermined, and the algorithm continues until the value of this maximum is met. The optimisation is carried out with a population size of 200 and 300 generations using the LHS for the initial population size in the final configuration of the research. For stability checking, the implementation has multi-seed runs, and records hypervolume progress as a convergence diagnostic.

Table 3-18 NSGA-II and LHS configuration used in the optimisation framework.

Item	Symbol	Value	Unit	Description	Role in model
		Optimisation framework			
Optimisation algorithm	—	NSGA-II	—	Elitist multi-objective evolutionary algorithm	Pareto search
Initialisation method	INIT_METHOD	LHS	—	Stratified sampling of the design space	Generation-zero population
Decision vector	(x)	$([R_p, R_t, OC, P_f])$	—	Design variables for panel rationalisation, truss rationalisation, opening/detailing complexity, and prefabrication level	Candidate encoding
Number of objectives	(M)	3	count	Cost, carbon, and process time	Multi-objective minimisation
Objective direction	—	Minimise all	—	Lower values are preferred for all objectives	Fitness definition
		Run configuration			
Population size	(N)	200	individuals	Number of candidate solutions per generation	Search resolution
Number of generations	(G)	300	generations	Maximum evolutionary iterations	Convergence control
Random seed, main run	—	42	—	Reproducibility setting for the main optimisation run	Baseline run
Multi-seed runs	—	11, 22, 33, 44, 55	—	Independent random seeds for repeated optimisation runs	Stability check
		Design-variable bounds			
Panel rationalisation	(R_p)	0.00–1.00	dimensionless	Panel standardisation intensity	Design constraint
Truss rationalisation	(R_t)	0.00–1.00	dimensionless	Truss standardisation intensity	Design constraint
Opening/detailing complexity	(OC)	0.80–1.30	dimensionless	Opening and detailing complexity factor	Design constraint
Prefabrication level	(P_f)	0.50–0.90	fraction	Off-site fabrication proportion	Design constraint
		Genetic operators			
Selection	—	NSGA-II environmental selection	—	Nondominated sorting with crowding-distance ranking	Diversity preservation
Crossover	—	Simulated binary, bounded	—	Recombination of continuous variables within bounds	Offspring generation

Item	Symbol	Value	Unit	Description	Role in model
Crossover distribution index	(η_c)	15	dimensionless	Controls the spread of crossover offspring around parent solutions	Exploration–exploitation balance
Mutation	—	Polynomial, bounded	—	Local perturbation of continuous variables within bounds	Search refinement
Mutation distribution index	(η_m)	20	dimensionless	Controls the spread of mutation perturbations	Search stability
Mutation probability	(p_m)	0.25	probability	Per-variable probability of mutation	Diversity control
Fitness and outputs					
Fitness vector	(F^*)	$([\bar{C}+P_C, \bar{E}+P_E, \bar{T}+P_T])$	dimensionless	Penalised normalised cost, carbon, and time objectives	Selection metric
Raw outputs	—	(C, E, T)	NZD, $kgCO_{2e}$, h	Unnormalised cost, carbon, and process-time outputs	Reporting and S2 comparison
Convergence and diagnostic outputs					
Hypervolume samples	—	500,000	samples	Monte Carlo estimation sample size for hypervolume calculation	Convergence tracking
Recording interval	—	Every 10 generations	generations	Frequency for recording optimisation progress	Progress assessment
LHS sensitivity samples	—	2,000	samples	Independent LHS evaluation set for design-space response analysis	Sensitivity analysis

Table 3-18 summarises the NSGA-II and LHS configuration used in the optimisation framework. LHS is used only to initialise the generation-zero population, while NSGA-II performs the iterative multi-objective search through nondominated sorting, crowding-distance ranking, crossover, mutation, and elitist environmental selection. The penalised normalised objective vector is used during selection, while the raw cost, carbon, and process-time values are retained for S2 comparison, interpretation, and dashboard reporting.

The crossover and mutation settings were selected as standard real-valued NSGA-II operator parameters for bounded continuous variables. The simulated binary crossover and polynomial mutation distribution indices provide controlled variation without treating the variables as discrete scheme selections. These settings are therefore appropriate for the continuous response vector x.

Although the response model contains four main decision variables, NSGA-II was retained as the primary search method because the implemented design space includes continuous variables, nonlinear interaction terms, feasibility screening, and post-processing of nondominated alternatives. The archive is not assumed to be globally exact; it is assessed through convergence behaviour, grid-reference comparison, multi-seed stability, feasibility checking, and S2-relative classification.

3.13: Model Verification, Practical Reasonableness, and Convergence Strategy

In this study, model credibility is assessed through verification, practical reasonableness checks, convergence diagnostics, sensitivity analysis, and robustness testing. These checks are used to determine whether the implemented response model behaves consistently, whether the optimisation search is stable, and whether selected Pareto solutions remain interpretable under the tested uncertainty ranges.

Five checks are used. First, algorithmic verification assesses the LHS-initialised NSGA-II output using hypervolume convergence, IGD, grid-reference comparison, feasibility counts, nondominated-solution counts, and multi-seed stability. Second, model-logic checking assesses whether the cost, carbon, and process-hour response functions follow the intended CFS framing mechanisms, including rationalisation, opening/detailing complexity, prefabrication level, logistics, waste, factory–site transfer, and lower-bound constraints. Third, practical reasonableness checking assesses whether the model-calculated S2 cost, A1–A5 carbon, and process-hour magnitudes are plausible when compared with the case-study area, steel mass, component breakdowns, coefficient audit, benchmark sources, and engineering judgement. Fourth, Monte Carlo robustness testing assesses selected representative solutions under selected coefficient ranges. Fifth, dashboard verification checks whether exported optimisation, sensitivity, robustness, ranking, and S2 comparison outputs are loaded and displayed consistently in the decision-support interface.

These checks support confidence in the framework as a comparative model-based decision-support method. They do not prove that the coefficients are physically exact, that the outputs are supplier quotations, that the carbon values are certified EPD or whole-building LCA results, that the solutions are structurally certified, or that the dashboard has been validated through formal industry user testing.

NSGA-II is an evolutionary algorithm based on randomness. Thus, a single Pareto plot is not sufficient for showing convergence. Hypervolume, grid-reference IGD, nondominated solution count, archive behaviour, and multi-seed convergence was used together in this study. The evaluation of convergence and diversity is to be done as a couple; NSGA-II itself uses nondominated sorting and crowding-distance preservation and it is important that both are evaluated (Deb et al., 2002).

In the validation process the calculation chain is identical. First, all candidate solutions are assessed for the cost, carbon and time objective functions. Second, the raw outputs are saved as C, E and T. Then, the values are normalised by the S2 denominators given in Section 3.9. Lastly, the objective space is normalised before computing the validation metrics. This makes validation directly tied to the same set of objective values that are used for optimisation.

The extent of the nondominated archive to dominate the normalised objective space was assessed through hypervolume, HV. In the current minimisation problem, a bigger HV usually means that the archive is shrinking to lower cost, lower carbon, and lower process time, but is still spread out. Considering both convergence and distribution, hypervolume is widely employed as a multi-objective performance indicator, but the interpretation of the indicator is dependent on the reference point used (Audet et al., 2021).

To ensure quality, a second check was performed using Grid-reference inverted generational distance (IGD). There is a need for a reference set for standard IGD. For this model case-study, a grid-reference approximation of the true Pareto front was employed due to the lack of this information. This is to prevent giving the impression that the true front is known, but at the same time providing a consistent external comparison set. The closer the nondominated archive is to the reference approximation the more accurate the corresponding IGD is. We have seen that IGD-type indicators are often used to compare a solution set acquired to a pre-specified reference point set in objective space (Ishibuchi et al., 2015).

Solution-count auditing was used to check whether the final archive contained enough usable design alternatives. The number of nondominated solutions, feasible solutions, and unique design vectors were recorded as N_{ND}, N_F, and N_U. These counts were not treated as objectives. They were diagnostic checks used to identify loss of diversity, excessive duplication, or premature convergence.

Multi-seed convergence was used to check the stability of the optimisation under different random seeds. Each seed used the same objective functions, variable bounds, constraints, LHS initialisation procedure, population size, number of generations, crossover, mutation, and validation metrics. Only the stochastic search pathway changed. A stable optimisation process should show similar final HV, similar final grid-reference IGD, and comparable N_{ND}, N_F, and $N_U$values across independent runs.

Table 3-19 Validation , reasonableness, and convergence metrics.

Item	Symbol	Unit	What it checks	How it is interpreted	Role in research
Normalised objective vector	$(\bar{F})$	dimensionless	Common comparison scale for cost, carbon, and time	Allows the three objectives to be assessed together	Basis for all validation metrics
Nondominated archive	(A_g)	solutions	Feasible nondominated solutions retained at generation (g)	A stable archive suggests consistent search behaviour	Main object used for convergence monitoring
Hypervolume	(HV)	dimensionless	Dominated objective-space volume relative to a fixed reference point	Higher (HV) indicates better convergence and spread	Tracks Pareto-front quality over generations
Grid-reference IGD	(IGD)	dimensionless	Distance from the archive to a grid-reference approximation	Lower (IGD) indicates closer agreement with the reference approximation	External convergence check
Nondominated solution count	(N_{ND})	count	Number of nondominated archive solutions	Very low values may indicate front collapse or poor diversity	Diversity diagnostic
Feasible solution count	(N_F)	count	Number of feasible candidate solutions	Low values may indicate overly restrictive constraints or many penalties	Feasibility diagnostic

Item	Symbol	Unit	What it checks	How it is interpreted	Role in research
Unique design-vector count	(N_U)	count	Number of non-duplicate design vectors	Low values may indicate duplication or search stagnation	Duplicate and diversity audit
Archive behaviour	—	—	Change in archive size and quality across generations	Stable growth or stabilisation supports convergence	Archive stability check
Multi-seed convergence	—	—	Consistency across independent random seeds	Similar (HV), (IGD), and count values indicate robust search behaviour	Stability check
Final convergence judgement	—	—	Combined reading of all validation indicators	No single metric is used alone	Supports reliability of final Pareto-front interpretation

Table 3-19 summarises the numerical indicators used for the algorithmic and diagnostic parts of this validation strategy, while the practical reasonableness and dashboard verification checks are reported in Chapter 4 and Chapter 7.

Hypervolume was calculated in normalised objective space using a fixed grid-derived reference box. The ideal point was defined as 0.995 times the minimum grid-evaluated normalised objective vector, while the reference point was defined as 1.050 times the maximum grid-evaluated normalised objective vector. For the reported run, this gave $\mathbf{z}_{HV}^{\min} = [0.854191, 0.931270, 0.732068]$ and $\mathbf{r}_{HV} = [1.259925, 1.091823, 1.180800]$. The same fixed reference box and sample set were used for all recorded generations, so changes in hypervolume reflect changes in the nondominated archive rather than changes in the reference point.

Because the hypervolume reference box was derived from the grid-evaluated objective range, the HV indicator is not treated as an external validation of the physical response model. Its role is narrower: it provides a fixed numerical reference for tracking whether the nondominated archive expands and stabilises during the NSGA-II run. The same fixed reference box was used for all generations and seeds, so the HV trend is suitable for convergence monitoring. However, physical validity is assessed separately through coefficient audit, S2 reasonableness checks, sensitivity analysis, Monte Carlo uncertainty propagation, and explicit limitation of the model boundary.

The validation and convergence strategy provides a numerical check on the reliability of the optimisation results. Hypervolume and grid-reference IGD assess the quality of the nondominated archive, while solution-count auditing and archive behaviour check whether the search retains sufficient diversity. Multi-seed convergence further tests whether the final optimisation output is stable under different random search pathways. Together, these indicators provide evidence of convergence stability, solution diversity, and reproducibility of the final nondominated archive.

3.14: Sensitivity Methodology

Section 3.14 defines the sensitivity analyses used to interpret the implemented response model. The purpose is not to validate the model empirically, but to identify which variables and coefficient groups most strongly affect CFS framing process cost, A1–A5 carbon, and process-hours within the adopted bounds. The sensitivity workflow includes one-at-a-time response testing around S2, LHS-based design-space screening, Spearman rank association, and PRCC-based partial association.

The sensitivity stage follows the same calculation chain used throughout the methodology:

objective functions → raw cost, carbon, and time values → S2-normalised values → sensitivity indicators

The raw outputs are C, E, and T, representing cost in NZD, carbon in $kgCO_{2e}$, and process time in hours. These are then normalised against S2 using the definitions in Section 3.9. The implemented framework also records the S2-relative percentage changes for each LHS design-space sample, which allows the sensitivity results to be interpreted against the same benchmark used in the optimisation results.

There were two levels of sensitivity analysis. First, one-at-a-time sensitivity was applied to this as a local diagnostic. One input is changed and the other inputs remain constant in the sensitivity OAT. This allows for easy interpretation of the results: a measure of the local effect of one variable around a specific reference condition. But interaction effects between variables are not fully reflected in OAT. This is why it is used only as a local screening tool, in line with the fact that single factor-at-a-time methods are not able to capture interactions in multi-input models (Saltelli & Annoni, 2010).

Second, model behaviour was tested using LHS design space screening to assess model behaviour over a broader range of input combinations. Latin Hypercube Sampling is a method that generates stratified samples of each range of input variables and is widely applied in computer experiments, where the objective is to sample a multidimensional design space uniformly as opposed to simple random sampling (McKay et al., 1979). The LHS sensitivity sample is not a part of LHS initialisation of NSGA-II in this study . The first use is for generation-zero optimisation population, the latter use is for design-space response screening.

The LHS screening dataset records the four optimiser variables R_p, R_t, OC, and P_f, the derived scheme labels, steel weight, raw cost, raw carbon, raw time, normalised outputs, S2 percentage changes, and feasibility-related outputs. This makes the sensitivity outputs traceable to the same objective functions and S2 comparison logic used in the optimisation framework.

The reporting categories panel_scheme and truss_scheme require additional care. These are derived categories from R_p and R_t, not independent optimiser-controlled variables. They are useful for interpretation, plots, and dashboard filtering, but their categorical nature means that PRCC values should not be over-interpreted as continuous sensitivity effects. The implemented framework therefore supplements rank-based sensitivity with grouped scheme-effect comparison for panel_scheme and truss_scheme. This reports count, mean, standard deviation, minimum, maximum, and mean difference from the global average for each category.

The sensitivity methodology is thus a mixture of local and screening-level evidence. An OAT is a basic, local test around selected reference conditions. Spearman and PRCC screening offer more information on input–output association throughout the design space, as compared to LHS. Grouped scheme-effect comparison offers a more secure interpretation of derived categorical reporting labels. These methods complement each other and can be used to interpret the influence of the variables without complete variance decomposition and without universal causal ranking.

The sensitivity calculations use the same objective-function outputs, S2 normalisation, design-variable bounds, and feasibility screening defined in earlier methodology sections. The LHS screening dataset records raw outputs, normalised outputs, S2-relative percentage changes, and feasibility indicators before Spearman, PRCC, and grouped scheme-effect summaries are calculated.

The sensitivity methodology provides a transparent check on which inputs most strongly influence the modelled cost, carbon, and process-time outputs. OAT sensitivity gives a local view of individual input effects, while LHS-based Spearman and PRCC screening provide broader evidence of input–output association across the design space. Derived scheme labels are interpreted through grouped comparisons to avoid overstating categorical variables as continuous effects. This approach supports later interpretation of variable influence while remaining consistent with the fixed CFS framing boundary and avoiding any variance-decomposition claim. Because R_p and R_t are continuous response-model variables, any grouped comparison by panel_scheme or truss_scheme is interpreted as reporting-level evidence only. The sensitivity results are interpreted as model-behaviour diagnostics within the implemented response formulation, not as empirical proof that the same variable hierarchy will apply to all CFS residential projects.

3.15: Monte Carlo Uncertainty Method

Section 3.15 defines the Monte Carlo robustness test applied to selected representative Pareto solutions. The conceptual role of Monte Carlo uncertainty propagation was established in Section 2.9. This section reports the implemented simulation design: selected uncertain coefficients, distribution assumptions, number of simulations, paired sampling against S2, percentile reporting, coefficient of variation, and probability of outperforming S2

The Monte Carlo method follows the same calculation chain used throughout the methodology:

selected Pareto solution → uncertain coefficient sampling → raw cost, carbon, and time values → paired S2 comparison → robustness indicators

Only Pareto representative solutions were used for the Monte Carlo analysis. These include cost centred, carbon centred, time centred and balanced alternatives. The same uncertain input samples were used for the evaluation of the S2 reference case. The paired comparison is significant because it implies that it is not just a comparison of an uncertain candidate solution with a predefined deterministic value for S2, but rather a comparison of a candidate solution to S2 under the same uncertainty sets. The implementation is based on this "paired logic" meaning that both candidate and S2 are evaluated with the same perturbed parameter set, and no S2 calibration scaling performed.

For each representative solution, the Monte Carlo output is reported to be using percentile values, coefficient of variation, and probability-based indicators. The $P5$, $P50$, and $P95$values describe the lower, median, and upper parts of the simulated output distribution. The coefficient of variation, CV, reports the relative spread of the output compared with its mean. The probability of outperforming S2 reports how often the selected solution produces lower cost, lower carbon, lower time, or all three objectives lower than S2 under paired uncertainty sampling. The implemented framework records $P5$, $P50$, $P95$, mean, standard deviation, CV, probability below the fixed S2 value, probability better than paired S2, and probability of dominating paired S2.

The main robustness indicators are interpreted as follows. A solution with a low median value but high CV may be attractive deterministically but less dependable under uncertainty. A solution with a slightly higher median but narrower uncertainty range and higher probability of outperforming S2 may be more robust for practical decision support. This is important in construction decision-making because the preferred framing option should not depend only on one deterministic cost, carbon, or time estimate.

Monte Carlo simulation is not treated as physical validation. It does not prove that the model reproduces real project cost, carbon, or site productivity. Measured project data from multiple CFS framing projects would be required for that level of validation. In this study, Monte Carlo analysis is used only to test whether the selected Pareto alternatives are stable under reasonable uncertainty in model coefficients.

The Monte Carlo analysis propagates uncertainty in selected coefficients and compares each representative solution with S2 under the same sampled coefficient set.

Table 3-20 Monte Carlo input ranges used for uncertainty propagation.

Item	Symbol	Range	Unit	Model use
Cost parameters				
Steel material rate	(c_{steel})	2.00–2.65	NZD/kg	Cost uncertainty
Factory processing rate	(c_{fac})	2.20–3.80	NZD/kg	Factory-cost uncertainty
Site labour rate	(c_{lab})	42.00–58.00	NZD/h	Labour-cost uncertainty
Transport cost rate	(c_{trans})	2.40–4.20	NZD/t-km	Freight-cost uncertainty
Panel setup cost	$(c_{p,\text{setup}})$	15.00–35.00	NZD/type	Panel setup-cost uncertainty
Truss setup cost	$(c_{t,\text{setup}})$	25.00–50.00	NZD/type	Truss setup-cost uncertainty
Carbon parameters				
Steel A1–A3 carbon factor	(EF_{steel})	1.50–2.89	$kgCO_{2e}$/kg	Product-stage carbon uncertainty
Freight carbon factor	(EF_{freight})	0.107–0.155	$kgCO_{2e}$/t-km	A4 transport-carbon uncertainty
A5 installation carbon factor	(EF_{A5})	1.50–4.50	$kgCO_{2e}/m^2$	Site-stage carbon uncertainty

Item	Symbol	Range	Unit	Model use
Time and interaction parameters				
Prefabrication time-saving coefficient	$\left(k_{P_f,T}\right)$	0.40–0.65	dimensionless	Site-time uncertainty
Prefabrication waste-reduction coefficient	$\left(k_{P_f,\text{waste}}\right)$	0.10–0.25	dimensionless	Waste and carbon uncertainty
Prefabrication A5 carbon-saving coefficient	$\left(k_{P_f,A5}\right)$	0.30–0.60	dimensionless	A5 carbon uncertainty
Logistics cost coefficient	$\left(k_{\log,C}\right)$	0.06–0.18	dimensionless	Logistics-cost uncertainty
Logistics time coefficient	$\left(k_{\log,T}\right)$	0.08–0.23	dimensionless	Logistics-time uncertainty
Logistics carbon coefficient	$\left(k_{\log,E}\right)$	0.02–0.07	dimensionless	Logistics-carbon uncertainty
Opening/detailing cost coefficient	$\left(k_{OC,C}\right)$	0.08–0.25	dimensionless	Detailing-cost uncertainty
Opening/detailing carbon coefficient	$\left(k_{OC,E}\right)$	0.05–0.14	dimensionless	Detailing-carbon uncertainty
Opening/detailing time coefficient	$\left(k_{OC,T}\right)$	0.12–0.35	dimensionless	Detailing-time uncertainty
Opening–prefabrication interaction factor	$\left(\lambda_{OC,P_f}\right)$	0.30–0.60	dimensionless	Interaction uncertainty
Weather downtime fraction	$\left(r_{\text{weather}}\right)$	0.05–0.15	fraction	Weather-time uncertainty
Weather shielding factor	$\left(s_{\text{weather}}\right)$	0.40–0.75	fraction	Weather-exposure uncertainty

Note: The Monte Carlo analysis uses $N_{MC} = 1000$ simulations and is interpreted as bounded robustness testing under selected coefficient ranges. It does not independently sample the fixed S2 base production times, does not test alternative response-model structures, and does not constitute Monte Carlo convergence testing. The sampled ranges are used to test selected representative Pareto solutions against S2 under consistent uncertainty assumptions, not to recalibrate the optimisation model or validate the outputs against measured project data.

The Monte Carlo input ranges in Table 3-20 represent selected uncertainty categories for cost, carbon, logistics, waste, prefabrication-carbon response, and weather-related assumptions. The resulting percentile ranges, coefficients of variation, robustness rankings, and probabilities of outperforming S2 are conditional on these selected ranges and on the implemented response-model structure. They should therefore be interpreted as model-based robustness evidence rather than contractor tender validation, product-specific environmental declaration evidence, measured productivity validation, or full model-form uncertainty analysis.

The Monte Carlo analysis propagates parameter uncertainty within the selected response-model structure. It does not test alternative model structures, such as non-linear rationalisation mappings, discrete fabricator-specific scheme libraries, different productivity equations, or alternative carbon-boundary definitions. These are treated as model-form limitations and are addressed through transparent disclosure, sensitivity interpretation, and future research recommendations. A full model-form sensitivity study, including alternative rationalisation knee points, prefabrication-response exponents, opening/detailing penalty structures, and correlated uncertainty assumptions, is outside the completed analysis and is identified as future work.

3.16: Dashboard Requirements and Evaluation

The dashboard is included in the methodology as the decision-support layer of the optimisation framework. The NSGA-II procedure produces a set of nondominated CFS framing alternatives rather than a single definitive answer. Therefore, an interface is needed to help users compare the alternatives, inspect trade-offs, review S2-relative performance, and understand robustness under uncertainty.

The dashboard is included in the methodology as the decision-support layer that receives exported optimisation and analysis outputs. Its purpose is to organise the final Pareto archive, S2 comparison values, representative solutions, sensitivity outputs, Monte Carlo robustness results, component breakdowns, ranking outputs, and scope notes into an interpretable decision-support workflow. The dashboard does not rerun NSGA-II, replace the imported Pareto front, perform structural verification, generate contractor pricing, or conduct whole-building LCA.

The dashboard uses exported optimisation outputs rather than recalculating the full optimisation inside the interface. The implementation exports a dashboard configuration file containing calibration values, Pareto solutions, representative solutions, sensitivity outputs, Monte Carlo results, and ranking outputs. The implemented framework writes these data to the dashboard data directory, including dashboard_config.json, which acts as the main data source

The dashboard is organised around the user tasks required for CFS framing decision support. These tasks include viewing the overall Pareto front, filtering solutions by cost, carbon, process time, prefabrication level, and robustness, comparing selected solutions against S2, reviewing uncertainty ranges, and ranking alternatives under different decision priorities. This aligns with the intended dashboard role in the research: compare, filter, rank, inspect uncertainty, export decision, and keep scope warnings visible.

The dashboard does not introduce new objective functions or change the optimisation results. It displays the outputs produced by the optimisation, validation, sensitivity, and Monte Carlo stages defined earlier in Chapter 3. Therefore, the dashboard is evaluated as a functional decision-support interface. It is not presented as certified structural design software, contractor tender validation, commercial estimating software, or evidence of industry adoption.

Table 3-21 Dashboard requirements and evaluation criteria.

Item	Requirement	Input / output	Evaluation criterion	Role in methodology
Data import	Load exported optimisation results	Pareto CSV / dashboard JSON	Dashboard loads results without manual editing	Links optimisation outputs to interface
S2 benchmark display	Show (C_{S2}), (E_{S2}), and (T_{S2})	S2 reference values	S2 values visible and clearly labelled	Baseline comparison
Pareto overview	Display cost–carbon–time trade-off solutions	Pareto-front solution set	User can inspect overall solution spread	Trade-off interpretation
Objective filtering	Filter by cost, carbon, time, and percentage change	Normalised and raw objective outputs	User can restrict solutions by decision priority	Practical screening

Item	Requirement	Input / output	Evaluation criterion	Role in methodology
Variable filtering	Filter by (R_p), (R_t), (OC), (P_f), panel scheme, and truss scheme	Design-variable and reporting-category data	User can inspect configuration effects	Design interpretation
Representative solutions	Display cost-focused, carbon-focused, time-focused, and balanced alternatives	Representative solution table	Key alternatives are identifiable	Practitioner decision support
S2 comparison	Compare selected solutions against S2	Raw and percentage-change values	Dashboard shows whether each solution improves or trades off against S2	Benchmark interpretation
Component breakdown	Show cost, carbon, and time component contributions	Component-level outputs	User can identify why a solution performs better or worse	Engineering explanation
Robustness display	Show Monte Carlo (P5), (P50), (P95), (CV), and probability better than S2	Monte Carlo output table	User can inspect uncertainty and robustness	Robust decision support
Sensitivity display	Show OAT/LHS sensitivity and grouped scheme effects	Sensitivity outputs	User can identify influential inputs and scheme effects	Model transparency
Ranking support	Rank alternatives by user-selected priorities	Weighted or ranking output table	User can compare alternatives under different preferences	Decision translation
Scope warning	Show model boundary and exclusions	Scope and validation notes	Dashboard states CFS framing only, A1–A5 only, process-hours only	Prevents overclaiming
Export support	Allow selected results to be exported or copied	Filtered solution table	User can document selected alternatives	Reporting support
Reproducibility	Use exported model outputs as official data source	Timestamped JSON / CSV outputs	Dashboard can be reloaded from saved outputs	Audit trail
Software status	State dashboard is not certified software	Methodology note / dashboard warning	Interface does not imply compliance, tender, or structural certification	Ethical and professional boundary

Table 3-21 defines the functional requirements used to evaluate the dashboard. The dashboard is evaluated by whether it can load the exported optimisation outputs, display the Pareto alternatives, compare solutions against S2, support filtering and ranking, show robustness and sensitivity information, and communicate the modelling boundary clearly. The evaluation is therefore based on decision-support functionality and traceability, not on commercial software certification or industry adoption. The dashboard evaluation is functional and traceability-based. It checks whether the interface displays and organises the exported model outputs correctly; it does not represent formal practitioner usability testing or industry validation.

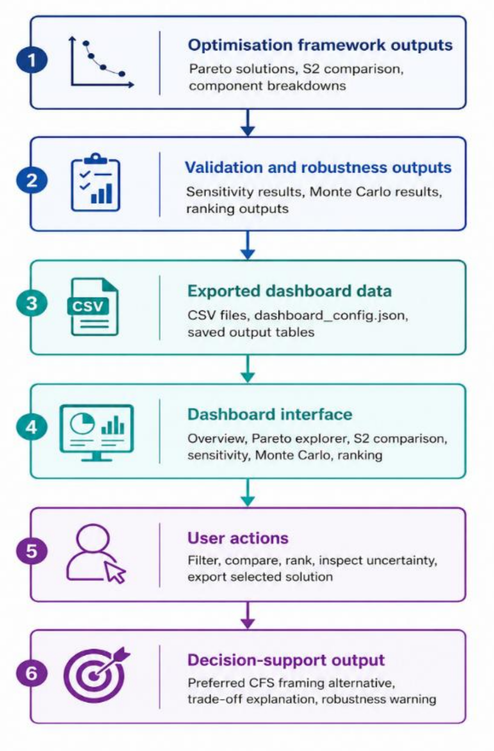

Figure 3-2 Dashboard data-flow architecture for exported model outputs.

The dashboard methodology defines how optimisation outputs are exported and organised for interpretation. It does not evaluate commercial usability or practitioner adoption. The dashboard is therefore treated as a decision-support translation layer that connects Pareto alternatives, S2-relative comparison, component breakdowns, sensitivity outputs, Monte Carlo robustness, and TOPSIS ranking within the fixed CFS framing case-study boundary.

Chapter 4: Model Implementation, Verification, and S2 Reasonableness Checks

4.1: Engineering Model Implementation Overview

Chapter 4 reports how the methodology defined in Chapter 3 was implemented and verified as an engineering response model for the S2 CFS framing case study. The purpose of the chapter is not to document the Python script line by line, but to show that the implemented model preserves traceability between the fixed case-study boundary, response coefficients, cost, carbon and process-hour objective functions, verification checks, diagnostic plots, and exported decision-support outputs.

The chapter is organised around engineering evidence. Section 4.2 confirms the active input data and fixed boundary conditions. Section 4.3 reports the calibration and assumption audit used for the implemented run. Sections 4.4–4.6 present the implemented cost, carbon, and process-hour component breakdowns. Section 4.7 examines material-weight and interaction effects. Section 4.8 consolidates the S2 verification, external reasonableness check, and diagnostic evidence used to judge whether the model-calculated S2 values are plausible within the adopted case-study boundary. Sections 4.9 and 4.10 then summarise, rather than document in detail, the export and reproducibility arrangements used to support later dashboard interpretation.

This chapter therefore functions as an implementation and verification bridge between the methodology in Chapter 3 and the optimisation results in Chapter 5. The key question is not how the code is organised, but whether the implemented response model reproduces the defined S2 reference, generates interpretable component breakdowns, and provides sufficient traceability for the subsequent optimisation, sensitivity, uncertainty, and decision-support analyses.

The optimisation stage applies the LHS-initialised NSGA-II configuration to explore the four-variable CFS framing decision space. The decision variables are panel rationalisation intensity, truss rationalisation intensity, opening/detailing complexity, and prefabrication level. The optimiser returns a nondominated archive, which is subsequently processed to extract the final interpreted Pareto front. This separation between the archive and the interpreted Pareto front is important because the archive records the optimisation search history, while the final Pareto front provides the decision set used for engineering interpretation in later chapters.

After optimisation, the framework performs several validation and post-processing stages. Grid-reference evaluation and grid-reference IGD are used to assess whether the NSGA-II approximation is consistent with a broader sampled reference set. Hypervolume convergence and nondominated-solution counts provide additional checks on optimisation quality and solution diversity. The workflow also selects representative configurations, computes signed percentage changes relative to S2, generates component-level breakdowns, and prepares ranked decision alternatives using TOPSIS. These outputs support the later results chapters by linking each recommended configuration to its design-variable settings, objective values, S2 comparison, and component-level cost, carbon, and time drivers.

The final stages of the pipeline prepare the analysis for interpretation and decision support. OAT and LHS-based sensitivity outputs are generated to identify influential variables and uncertain coefficients, while Monte Carlo uncertainty analysis tests whether selected configurations remain preferable under plausible variation in cost, carbon, and productivity assumptions.

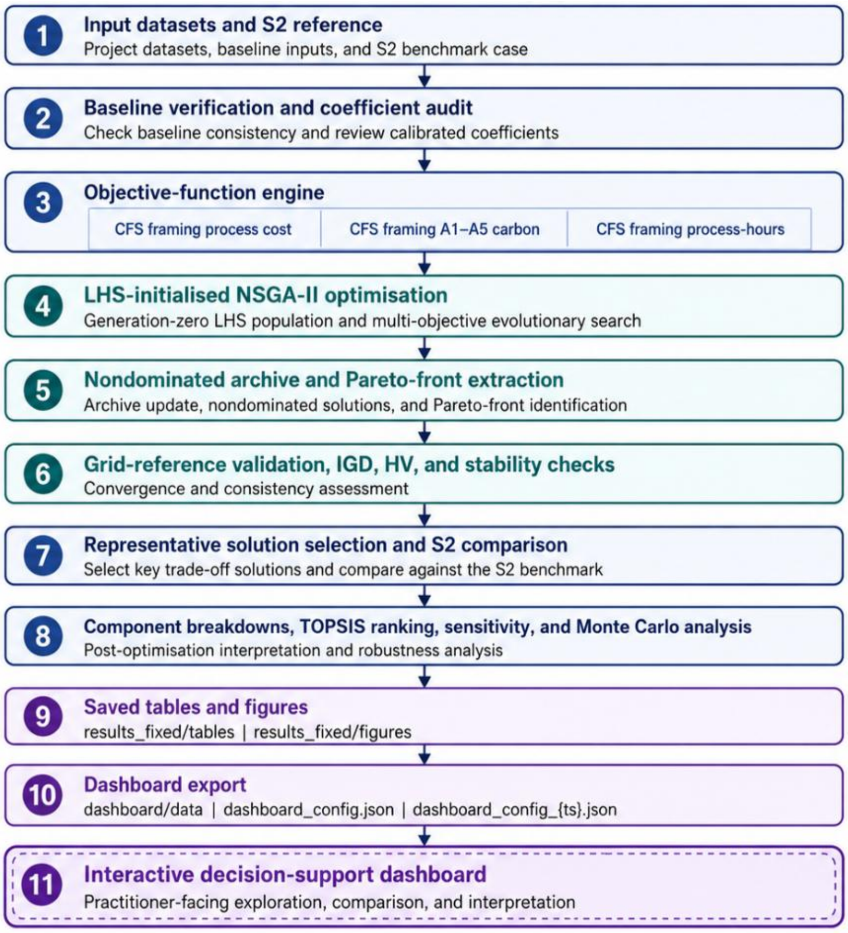

Figure 4-1 Implemented model traceability workflow for the CFS framing response model.

The figure shows the traceability pathway from fixed case-study inputs and the model-calculated S2 reference through objective-function evaluation, verification checks, optimisation outputs, sensitivity and robustness outputs, and dashboard-ready export files. The figure is included to show model-output traceability, not software architecture or code structure.

Table 4-1 Engineering implementation traceability summary for the reported model run.

Implementation stage	Engineering purpose	Main evidence retained in thesis
Fixed input and S2 reference loading	Confirms the case-study boundary used for all comparisons	Table 4-2, Table 4-3
Calibration and coefficient audit	Confirms cost, carbon, time, and interaction assumptions used in the implemented run	Table 4-4 and Chapter 3 coefficient audit tables
Cost, carbon, and process-hour evaluation	Confirms that the implemented objectives reproduce the defined response model	Tables 4-5, 4-6, 4-7
Material-weight and interaction diagnostics	Confirms that rationalisation, opening/detailing, and prefabrication effects behave as intended	Figures 4-5 and 4-6
S2 verification and reasonableness checking	Confirms the model-calculated S2 values and compares them with plausibility references	Table 4-8
Optimisation-output export	Confirms that the verified model produces traceable outputs for Chapters 5–7	Appendix output archive table

Table 4-1 summarises the implementation evidence needed for engineering interpretation. The detailed file-path and timestamped-output inventory is retained in Appendix 10 because it supports reproducibility but does not need to dominate the main chapter.

This output-driven structure is particularly important for the present research because the model is used to support engineering interpretation rather than merely to demonstrate algorithmic execution. The saved files link each CFS framing configuration to its design-variable settings, objective values, S2 percentage comparison, component breakdowns, validation metrics, ranking position, and uncertainty performance. As a result, the framework provides a transparent bridge between computational optimisation and civil/structural engineering decision support.

4.2: Input Data and Fixed Boundary Conditions

The implemented framework was developed around a fixed CFS residential framing case-study boundary. This boundary defines the geometry, framing quantities, baseline steel mass, and S2 reference objectives used for optimisation, comparison, and dashboard interpretation. The purpose of this stage is to ensure that the optimisation explores configuration decisions within a controlled civil/structural engineering case, rather than changing the building geometry or structural system during the search process.

The reported model run was defined using the final case-study boundary and response-model parameters documented in Chapter 3. Earlier project files and draft data tables were retained only as background records and were not treated as active inputs for the reported optimisation. This avoids ambiguity between preliminary data extraction work and the final model configuration used to generate the study results.

The active S2 reference is the project-derived, model-calculated configuration defined in Section 3.2. Its geometry, wall-panel count, roof-truss count, member information, and steel mass are derived from the project framing information, while its cost, carbon, and process-hour values are calculated through the final response model.

Table 4-2 Input data inventory used by the implemented CFS framing optimisation framework.

Input or configuration source	Active value source	Main content	Role in implementation	Output evidence
Project geometry and framing quantities	Current model configuration	Footprint area, envelope area, wall-panel count, roof-truss count, and baseline steel mass	Defines the fixed CFS residential framing case-study boundary	active_input_inventory_{ts}.csv; fixed_boundary_conditions_{ts}.csv
S2 reference design vector	Current model configuration	$(R_p = 0.0), (R_t = 0.0), (OC = 1.00), (P_f = 0.72)$	Defines the model-calculated S2 benchmark configuration	s2_reference_values_{ts}.csv
Cost coefficients	Current model configuration	Steel material rate, factory processing rate, site labour rate, setup cost, transport cost, waste handling, and overhead	Supports CFS framing process cost calculation	calibration_audit_{ts}.csv
Carbon coefficients	Current model configuration	A1–A3 steel carbon factor, A4 freight factor, A5 installation factor, factory energy factor, and waste carbon allowance	Supports CFS framing A1–A5 carbon calculation	calibration_audit_{ts}.csv
Time parameters	Current model configuration	Factory production time, setup time, delivery/logistics time, site mobilisation, installation time, and weather allowance	Supports CFS framing process-hours calculation	calibration_audit_{ts}.csv; component_breakdown_time_{ts}.csv
Decision-variable bounds	Current model configuration	Ranges for (R_p), (R_t), (OC), and (P_f)	Defines the LHS and NSGA-II search space	design_variable_table_{ts}.csv; initial_population_lhs_audit_{ts}.csv
Panel and truss reporting mappings	Current objective-function engine	Derived panel_scheme and truss_scheme labels from continuous (R_p) and (R_t)	Supports tables, plots, and dashboard filtering	pareto_front_{ts}.csv
Feasibility and constructability limits	Current model configuration	Panel width, panel height, stud spacing, steel thickness, lift mass, factory repetition score, and opening complexity threshold	Screens candidate configurations for practical feasibility	validation_status_{ts}.csv; pareto_front_{ts}.csv
Dashboard export files	Current output pipeline	Dashboard-ready CSV files and configuration JSON files	Transfers optimisation outputs to the decision-support interface	dashboard/data/*.csv; dashboard_config. json; dashboard_config_{ts}. json

Table notes: ts} denotes the timestamp associated with a specific model run. The exact timestamped filename should match the saved output produced during the reported run.

Table 4-3 Fixed boundary conditions and S2 reference values used for optimisation and comparison.

Boundary item	**Value**	**Unit**	**Status in model**	**Output evidence**
Case-study type	Single-storey CFS residential framing case	—	Fixed	fixed_boundary_conditions.csv
Footprint area	98.292	m²	Fixed	fixed_boundary_conditions.csv
Envelope area	217.2555	m²	Fixed	fixed_boundary_conditions.csv
Total wall panels	48	count	Fixed	fixed_boundary_conditions.csv
Total roof trusses	58	count	Fixed	fixed_boundary_conditions.csv
Baseline steel mass	2195.264	kg	Fixed reference quantity	fixed_boundary_conditions.csv
CFS section source	F325iT – Imported Section	—	Fixed	fixed_boundary_conditions.csv
CFS section designation	89 S 41	—	Fixed	fixed_boundary_conditions.csv
Steel thickness	0.75	mm	Fixed; not optimised	fixed_boundary_conditions.csv
Steel grade	G500	—	Fixed	fixed_boundary_conditions.csv
Coating class	Z275	—	Fixed	fixed_boundary_conditions.csv
Panel height	2.7	m	Fixed project boundary condition	fixed_boundary_conditions.csv
Stud spacing	600.0	mm	Fixed project boundary condition	fixed_boundary_conditions.csv
Maximum panel width	2.4	m	Hard feasibility constraint	fixed_boundary_conditions.csv
Maximum lift mass	2000.0	kg	Hard feasibility constraint	fixed_boundary_conditions.csv
S2 reference vector	([0.0, 0.0, 1.0, 0.72])	dimensionless / fraction	Reference design vector	s2_reference_values.csv
S2 CFS framing process cost	18,335.47	NZD	Model-calculated reference objective	s2_reference_values.csv
S2 CFS framing A1–A5 carbon	5,366.79	$kgCO_{2e}$	Model-calculated reference objective	s2_reference_values.csv
S2 CFS framing process-hours	93.92	h	Model-calculated reference objective	s2_reference_values.csv
Optimised building geometry	Not varied	—	Excluded from optimisation	validation_status.csv
Structural member sizing	Not varied	—	Excluded from optimisation	validation_status.csv
Whole-building LCA scope	Not included	—	Outside model boundary	validation_status.csv
Contractor tender benchmarking	Not included	—	Outside model boundary	validation_status.csv

Note: Table 4-3 records the fixed boundary used in the reported computational run. The S2 cost, carbon, and process-hour values are model-calculated reference values used for normalisation and comparison.

This structure supports reproducibility because a second researcher can inspect the exported configuration tables, verify the fixed boundary conditions, and rerun the optimisation against the same S2 reference. It also protects the engineering interpretation of the results: changes in CFS framing process cost, CFS framing A1–A5 carbon, and CFS framing process-hours are attributed to configuration decisions rather than uncontrolled changes in building size, framing quantity, or structural system.

4.3: Calibration Audit and Assumption Register

The calibration audit records how the final response model was anchored to the S2 reference case. The audit separates fixed case-study data, such as geometry, panel count, truss count, section designation, and steel mass, from model-calculated objective values and scenario-based coefficients. This separation is important because the implementation uses S2 as a consistent reference for comparison, not as a measured commercial benchmark or hidden optimisation constraint.

The implemented framework includes a calibration audit stage to make the model assumptions transparent before optimisation and decision-support outputs are interpreted. This stage records the coefficient values used in the cost, carbon, time, sensitivity, and Monte Carlo components of the model. The purpose is not only to document the numerical inputs, but also to identify the source basis, validation status, uncertainty range, and model use associated with each parameter. This is important because the framework is applied to a fixed CFS residential framing case study under New Zealand conditions, where several construction rates, carbon factors, and productivity assumptions must be treated as bounded estimates rather than exact physical measurements.

The calibration audit is exported as calibration_audit.csv. This output provides the main evidence that the implemented model can be reviewed and recalibrated by another researcher. Each coefficient is recorded with its value, unit, range, source type, confidence level, and model use. The source type distinguishes whether the value is directly measured from the case-study dataset, derived from model calculations, estimated from market or project assumptions, calibrated to preserve consistency with the S2 reference case, or adopted from literature and published databases. This classification prevents the model from presenting all parameters with the same evidential status.

The assumption register complements the calibration audit by summarising the fixed case-study conditions, reference assumptions, and uncertain coefficients used in the reported model run.

The calibration audit separates fixed case-study boundary assumptions from uncertain response coefficients. Fixed boundary assumptions include the selected CFS section, case-study geometry, S2 reference mass, and the process structure used to define the S2 benchmark. Scenario-based assumptions include transport distance, freight carbon factor, weather downtime, and weather shielding. These assumptions are retained to keep the model reproducible and bounded, while uncertainty is tested through the parameters selected for sensitivity and Monte Carlo analysis. This distinction is important because the framework is intended for early-stage CFS framing decision support.

Table 4-4 Calibration audit for the reported model run.

Parameter	Value	Unit	Range	Source basis	Validation status	Model use
C_S2_reference	18,335.47	NZD	Not applicable	Model-calculated CFS framing process-cost reference	Calculated from current objective-function engine	S2 benchmark for normalisation and comparison
CO2_S2_reference	5,366.79	$kgCO_{2e}$	Not applicable	Model-calculated CFS framing A1–A5 carbon reference	Calculated from current objective-function engine	S2 benchmark for normalisation and comparison
Time_S2_reference	93.92	process-hours	Not applicable	Model-calculated CFS framing process-hour reference	Calculated from current objective-function engine	S2 benchmark for normalisation and comparison
footprint_area_m2	98.292	m²	Not applicable	Declared model boundary case-study geometry	Model configuration value	Fixed CFS framing boundary
envelope_area_m2	217.2555	m²	Not applicable	Declared model boundary envelope/framing area	Model configuration value	A5 carbon and reporting normalisation
num_wall_panels	48	count	Not applicable	Declared model boundary case-study panel quantity	Model configuration value	Panel rationalisation and constructability metrics
num_roof_trusses	58	count	Not applicable	Declared model boundary case-study truss quantity	Model configuration value	Truss rationalisation and constructability metrics
W_base	2195.264	kg	Not applicable	Declared model boundary baseline CFS steel mass	Model configuration value	Material cost and A1–A3 steel carbon
raw_steel_material_nzd_per_kg	2.40	NZD/kg	2.00–2.65	Declared model boundary market-rate assumption	Uncertainty-tested parameter	CFS framing process cost
factory_processing_nzd_per_kg	2.80	NZD/kg	2.20–3.80	Declared model boundary factory processing assumption	Uncertainty-tested parameter	CFS framing process cost

Parameter	Value	Unit	Range	Source basis	Validation status	Model use
site_labor_rate_nzd_per_hr	48.00	NZD/h	42.00–58.00	Declared model boundary labour-rate assumption	Uncertainty-tested parameter	CFS framing process cost and process-hours coupling
transport_cost_nzd_per_tonne_km	3.00	NZD/t-km	2.40–4.20	Declared model boundary freight cost assumption	Uncertainty-tested parameter	Transport cost
steel_carbon_factor_a1_a3_kg CO_{2e}_per_kg	2.10	$kgCO_{2e}$/kg	1.50–2.89	Declared model boundary A1–A3 steel carbon factor	Uncertainty-tested parameter	CFS framing A1–A5 carbon
freight_$kgCO_{2e}$_per_tonne_km	0.135	$kgCO_{2e}$/t-km	0.107–0.155	Declared model boundary A4 freight carbon factor	Uncertainty-tested parameter	A4 transport carbon
a5_installation_carbon_$kgCO_{2e}$_per_m2	3.20	$kgCO_{2e}/m^2$	1.50–4.50	Declared model boundary A5 installation carbon factor	Uncertainty-tested parameter	A5 installation carbon
T_factory_truss_base	24.20	h	19.36–29.04	Case-study S2 process-hour assumption, bounded by panelised/off-site production literature and project workflow logic	Model configuration value	Fixed S2 truss factory-production reference
T_factory_panel_base	11.70	h	9.36–14.04	Case-study S2 process-hour assumption, bounded by panelised/off-site production literature and project workflow logic	Model configuration value	Fixed S2 panel factory-production reference
T_installation_truss_base	18.18	h	Not specified	Declared model boundary process-hour assumption	Model configuration value	CFS framing process-hours
T_installation_panel_base	57.06	h	Not specified	Declared model boundary process-hour assumption	Model configuration value	CFS framing process-hours
k_prefab_site_time_saving	0.60	dimensionless	0.40–0.65	Declared model boundary prefabrication response coefficient	Uncertainty-tested response coefficient	CFS framing process-hours
k_prefab_a5_carbon_saving	0.35	dimensionless	0.30–0.60	Declared model boundary prefabrication carbon response coefficient	Uncertainty-tested response coefficient	CFS framing A1–A5 carbon
k_logistics_cost	0.11	dimensionless	0.06–0.18	Declared model boundary logistics response coefficient	Uncertainty-tested response coefficient	CFS framing process cost
k_logistics_time	0.15	dimensionless	0.08–0.23	Declared model boundary logistics response coefficient	Uncertainty-tested response coefficient	CFS framing process-hours

Parameter	Value	Unit	Range	Source basis	Validation status	Model use
k_logistics_carbon	0.04	dimensionless	0.02–0.07	Declared model boundary logistics carbon response coefficient	Uncertainty-tested response coefficient	CFS framing A1–A5 carbon
$\boldsymbol{OC}$_range	0.80–1.30	dimensionless	0.80–1.30	Declared model boundary decision-variable range	Optimisation search bound	Opening/detailing complexity variable
$\boldsymbol{P_f}$_range	0.50–0.90	fraction	0.50–0.90	Declared model boundary decision-variable range	Optimisation search bound	Prefabrication-level variable
$\boldsymbol{R_p}$_range	0.00–1.00	dimensionless	0.00–1.00	Declared model boundary decision-variable range	Optimisation search bound	Panel rationalisation variable
$\boldsymbol{R_t}$_range	0.00–1.00	dimensionless	0.00–1.00	Declared model boundary decision-variable range	Optimisation search bound	Truss rationalisation variable

The calibration audit is used as an assumption-validity tool, not only as a reproducibility output. Its purpose is to distinguish project-derived quantities, model-calculated reference values, literature-supported coefficients, scenario assumptions, calibrated response coefficients, and uncertainty-tested parameters before optimisation results are interpreted. This distinction is central to model credibility because the S2 geometry, panel count, truss count, and steel mass have stronger evidential support than several productivity, logistics, waste, and interaction coefficients. The audit therefore identifies which parts of the model are directly project-derived and which parts require sensitivity testing, Monte Carlo uncertainty propagation, or future recalibration using measured CFS project data.

4.4: Cost Model Implementation

The cost implementation translates the cost objective defined in Chapter 3 into component-level outputs for the reported model run. The purpose of this section is to verify that the implemented response model produces the S2 cost benchmark and component breakdown used later for normalisation, S2-relative comparison, and Pareto interpretation. The implementation is therefore discussed through component traceability rather than through a repeated derivation of the full cost equation.

The cost model evaluates each candidate configuration using the four implemented design variables: panel rationalisation intensity R_p, truss rationalisation intensity R_t opening/detailing complexity OC, and prefabrication level P_f. These variables are first converted into an engineering state that includes effective panel-class and truss-type counts, repetition indices, opening/detailing effects, prefabrication response, overload effects, and logistics sensitivity. This intermediate engineering state allows the cost model to represent the physical consequences of configuration decisions rather than applying a single generic multiplier.

The implemented cost model returns both the total cost objective and its component-level breakdown. The main cost components are steel material, factory processing, setup and jigging, site installation labour, transport, delivery logistics, site mobilisation, opening/detailing standardisation, opening/detailing complexity, waste handling, and overhead. This structure is important because later optimisation results can be interpreted through the construction mechanisms that cause the cost change. For example, a solution may reduce site labour through higher prefabrication but increase factory processing or logistics cost. Similarly, rationalisation may reduce setup effort while introducing grouping, handling, or oversizing penalties.

Panel and truss rationalisation are therefore implemented as balanced response effects rather than as unconditional savings. Moderate rationalisation can improve repetition and reduce setup effort, while excessive rationalisation can increase grouping and logistics penalties. Opening/detailing complexity is also treated as an active design-for-manufacture response: lower values represent more standardised lintel, jamb, trimming, and service-penetration details, while higher values represent custom detailing and additional coordination. Prefabrication level affects the factory–site cost balance by modifying factory processing effort, site labour, logistics sensitivity, and overload-related penalties. This enables optimisation to evaluate realistic cost trade-offs across fabrication, transport, and site installation stages.

The cost implementation is verified through the cost component breakdown output generated during the reported model run. This output provides the link between the mathematical cost formulation in Section 3.5 and the cost values used in later Pareto-front interpretation. It also supports reproducibility because another researcher can inspect the reported component terms and confirm how the total cost value was assembled for representative configurations.

Table 4-5 S2 cost component breakdown from the implemented cost model.

Cost component	**Model output value**	**Unit**	**Share of S2 total cost**	**Implementation interpretation**
Steel material	5,268.63	NZD	28.73%	Material supply contribution based on the estimated CFS steel mass
Factory processing	6,822.88	NZD	37.21%	Factory production, roll-forming, cutting, handling, and processing effort
Setup and jigging	1,495.00	NZD	8.15%	Setup cost associated with panel and truss families
Site installation labour	1,989.40	NZD	10.85%	Labour cost associated with CFS framing installation on site
Transport	823.18	NZD	4.49%	Freight-related cost for transporting CFS framing elements
Delivery logistics	98.08	NZD	0.53%	Delivery coordination, scheduling, and unloading interface cost
Site mobilisation	144.00	NZD	0.79%	Site-start, access, safety, and mobilisation allowance
Opening/detailing standardisation	0.00	NZD	0.00%	No standardisation adjustment at the S2 reference opening complexity
Opening/detailing complexity	0.00	NZD	0.00%	No complexity penalty at the S2 reference opening complexity
Waste handling	27.44	NZD	0.15%	Steel waste/offcut handling allowance
Overhead	1,666.86	NZD	9.09%	Project-level indirect and coordination allowance
Total CFS framing process cost	18,335.47	NZD	100.00%	Model-calculated S2 cost reference

Note: This component breakdown corresponds to the same model-calculated S2 reference used in Table 3.2. All S2 cost components are reported from the final implemented response model at $\mathbf{x}_{S2} = [0.00, 0.00, 1.00, 0.72]$.

As an absolute-value reasonableness check, the model-calculated S2 CFS framing process cost can be normalised by the case-study envelope/framing area:

$$\frac{C_{S2}}{A_e} = \frac{18335.47}{217.2555} = 84.39\ \mathrm{NZD/m^2}$$

$$\frac{C_{S2}}{A_f} = \frac{18335.47}{98.292} = 186.54\ \mathrm{NZD/m^2}$$

These values are presented only as a reasonableness check of the implemented S2 reference cost and should not be interpreted as a market quotation. SteelHaus publicly states that steel-frame rates can range from \$70 + GST to \$330 + GST per m^2, with a typical average of \$160 + GST to \$200 + GST per m^2, depending on scope and complexity. Use it only as external context, not as formal validation (*SteelHaus*, 2024).

For external context, publicly available New Zealand light-gauge steel framing price guidance indicates that quoted steel-frame rates can vary widely depending on scope, project complexity, height, levels, and project-specific conditions. The S2 value is therefore used as a model-calculated benchmark for relative optimisation and S2-normalised comparison, not as an industry price claim or a replacement for supplier quotation.

Table 4-5 shows that S2 cost reference is dominated by factory processing, steel material, site installation labour, overhead, and setup/jigging costs. This confirms that the implemented cost objective is not controlled by a single rate or lump-sum allowance. Instead, the total CFS framing process cost is assembled from traceable material, factory, site, logistics, waste, and overhead components. This component structure is important for later Pareto-front interpretation because cost changes can be linked back to fabrication, installation, rationalisation, and prefabrication mechanisms.

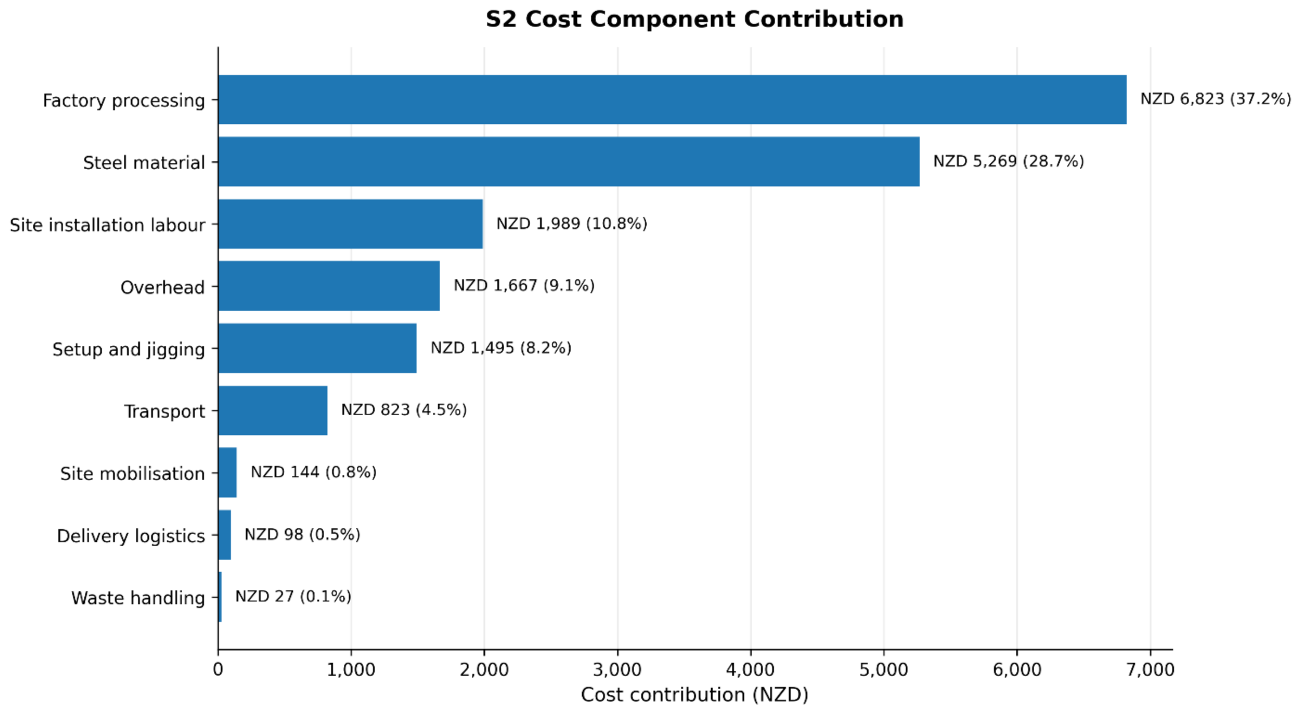

Figure 4-2 S2 cost component contribution from the implemented CFS framing process-cost model.

The figure shows how the model-calculated S2 cost is assembled from material, factory, setup, site labour, logistics, waste, and overhead components. (Zero-value opening/detailing terms are retained in the table but omitted from the figure.)

4.5: Carbon Model Implementation

The carbon formulation defined in Section 3.6 was implemented as the second objective function. The purpose of this section is to show implementation traceability rather than restate the full carbon equation. The implemented carbon model reports CFS framing A1–A5 carbon through separate A1–A3 steel, waste/offcut, A4 transport, factory energy, and A5 installation components. These outputs allow the model-calculated S2 carbon value and Pareto-solution carbon differences to be traced to their component sources.

The implementation applies the four design variables, R_p, R_t, OC, and P_f, through the same engineering response structure used across the optimisation framework. Panel and truss rationalisation influence material-weight response, waste allowance, transport sensitivity, and grouping effects. Opening/detailing complexity affects local detailing, trimming, coordination, and installation-stage carbon response. Prefabrication level influences factory production energy, A5 site installation carbon, logistics sensitivity, and waste reduction. This structure allows the carbon objective to respond to realistic CFS framing configuration changes rather than being treated as a fixed steel-mass multiplier only.

The implementation also avoids double counting site-stage energy. A5 installation carbon is treated as the site-stage construction carbon term, while factory energy is calculated separately from factory production time and electricity demand. Site energy is therefore not added again as a separate component where it would overlap with the A5 installation allowance. This distinction is important for maintaining a consistent A1–A5 boundary.

The carbon component breakdown output verifies the implementation by separating the S2 reference carbon into A1–A3 steel, waste/offcut carbon, A4 transport, factory energy, and A5 installation components. The S2 breakdown confirms that steel production dominates the CFS framing A1–A5 result, while A5 installation forms the main construction-stage contribution. Carbon intensity is reported separately as $kgCO_{2e}/m^2$ of floor area for interpretation only and should not be read as a whole-building carbon-intensity result.

The carbon coefficients used in the implementation are documented in the calibration audit; therefore, this section does not repeat coefficient justification. The purpose of this section is to show that the Chapter 3 carbon formulation was implemented as a traceable component-level model and that the exported carbon breakdown can be used to verify the reported objective values.

Table 4-6 S2 carbon component breakdown from the implemented carbon model.

Carbon component	**Model output value**	**Unit**	**Share of S2 total carbon**	**Implementation interpretation**
A1–A3 steel	4,610.05	$kgCO_{2e}$	85.90%	Embodied carbon associated with CFS steel production and supply
Waste/offcut carbon	202.92	$kgCO_{2e}$	3.78%	Carbon allowance associated with steel waste and offcuts
A4 transport	36.14	$kgCO_{2e}$	0.67%	Freight-related carbon for transporting CFS framing elements
Factory energy	16.57	$kgCO_{2e}$	0.31%	Electricity-related carbon from factory production activity
A5 site installation	501.11	$kgCO_{2e}$	9.34%	Installation-stage carbon associated with the CFS framing process
Total CFS framing A1–A5 carbon	5,366.79	$kgCO_{2e}$	100.00%	Model-calculated S2 carbon reference

Factory energy is included as an explicit but minor carbon term to avoid hiding production-stage electricity use inside a single global carbon factor. The reported model uses $P_{\text{fac}} = 5.0\text{kW}$ and $EF_{\text{grid}} = 0.08$ $kgCO_{2e}$/kWh. At the S2 reference condition this gives $E_{\text{fac,S2}} = 14.36$$kgCO_{2e}$, which is approximately 0.27% of the total S2 CFS framing A1–A5 carbon. A plausibility range of 3–10 kW is therefore reported for transparency, but the term is not a dominant driver of the carbon objective. Even doubling $P_{\text{fac}}$from 5 kW to 10 kW would increase S2 carbon by approximately 14.36 $kgCO_{2e}$, or about 0.27% of the S2 total. The factory power assumption is therefore disclosed as a low-impact scenario coefficient rather than treated as a primary source of optimisation uncertainty.

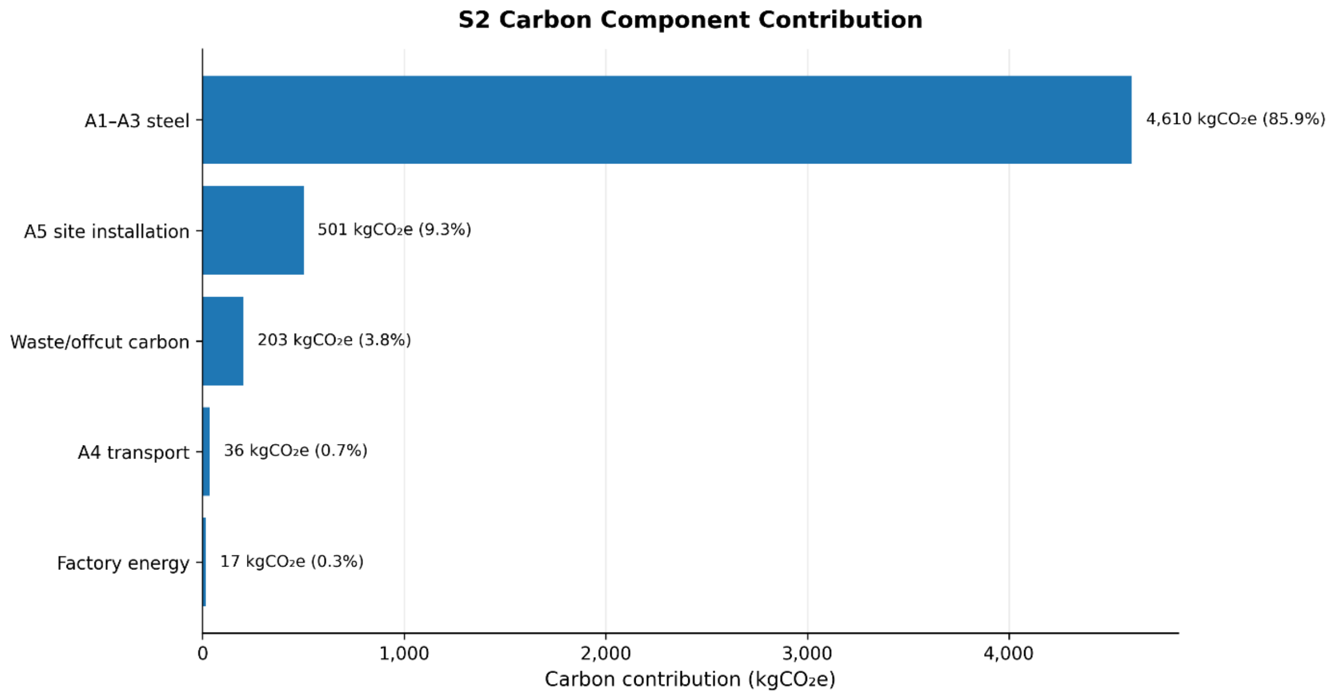

Figure 4-3 S2 carbon component contribution from the implemented CFS framing A1–A5 carbon model. The figure shows the relative contribution of A1–A3 steel, A4 transport, factory energy, A5 installation, and waste/offcut carbon within the CFS framing boundary.

4.6: Time Model Implementation

The time implementation reports CFS framing process-hours rather than calendar construction duration. The implemented response includes factory production, setup, delivery/logistics, site installation, site mobilisation, and weather allowance, consistent with the scope defined in Chapter 3. The purpose of this section is to verify how the process-hour objective was assembled for S2 and how the output remains distinct from a CPM schedule or full project programme.

The implemented time model separates the total process-hours into factory production, factory setup, delivery logistics, site mobilisation, site installation, and weather allowance. This component structure reflects the practical distinction between off-site fabrication and onsite assembly in prefabricated CFS framing. Factory production captures roll-forming, cutting, handling, and fabrication activity. Factory setup represents jigging, CNC preparation, and production changeover. Delivery logistics represents transport coordination and unloading interface time. Site mobilisation captures site-start and access preparation activities, while site installation represents the onsite placement and fixing of wall panels and roof trusses. A weather allowance is included to represent site productivity risk within the New Zealand case-study context.

The four design variables influence the time objective through the same engineering response structure used in the cost and carbon models. Panel rationalisation R_p and truss rationalisation R_t can reduce repetition-related setup and site handling effort, but excessive rationalisation may introduce grouping, handling, and logistics penalties. Opening/detailing complexity OCaffects site installation and factory coordination because non-standard lintels, jamb details, trimming, and service penetrations increase fabrication and installation difficulty. Prefabrication level $P_f$affects the factory–site balance: higher prefabrication can reduce site installation time, but it also increases factory production effort and may introduce logistics or factory overload effects at high prefabrication levels.

This implementation structure is important because it prevents prefabrication from being treated as a simple one-directional time saving. Instead, the model represents the trade-off between increased factory effort and reduced site effort. The S2 time component breakdown shows that the reference process-hours are distributed between factory production and site installation, which supports the decision to model time as a factory–site process balance rather than as a single productivity multiplier.

The time component breakdown output verifies the implementation by separating the S2 reference time into factory production, factory setup, delivery logistics, site mobilisation, site installation, and weather allowance. This output provides traceability between the Chapter 3-time formulation and the time values used in optimisation, S2 comparison, Pareto-front interpretation, and dashboard reporting.

Table 4-7 S2 time component breakdown from the implemented time model.

Time component	**Model output value**	**Unit**	**Share of S2 total time**	**Implementation interpretation**
Factory production	41.43	Hr	44.11%	Factory fabrication, roll-forming, cutting, and production time
Factory setup	3.50	Hr	3.73%	Jig setup, CNC preparation, and production changeover allowance
Delivery logistics	2.06	Hr	2.19%	Delivery coordination and unloading interface time
Site mobilisation	3.00	hr	3.19%	Site setup, access preparation, and mobilisation time
Site installation	41.45	Hr	44.13%	On-site installation of CFS wall panels and roof trusses
Weather allowance	2.49	Hr	2.65%	Waikato weather-related process-time allowance
Total CFS framing process-hours	93.92	Hr	100.00%	Model-calculated S2 time reference

Table 4-7 shows that S2 process time is evenly divided between factory production and site installation. This supports the modelling decision to treat prefabrication level as a factory–site trade-off rather than as a simple time-saving multiplier.

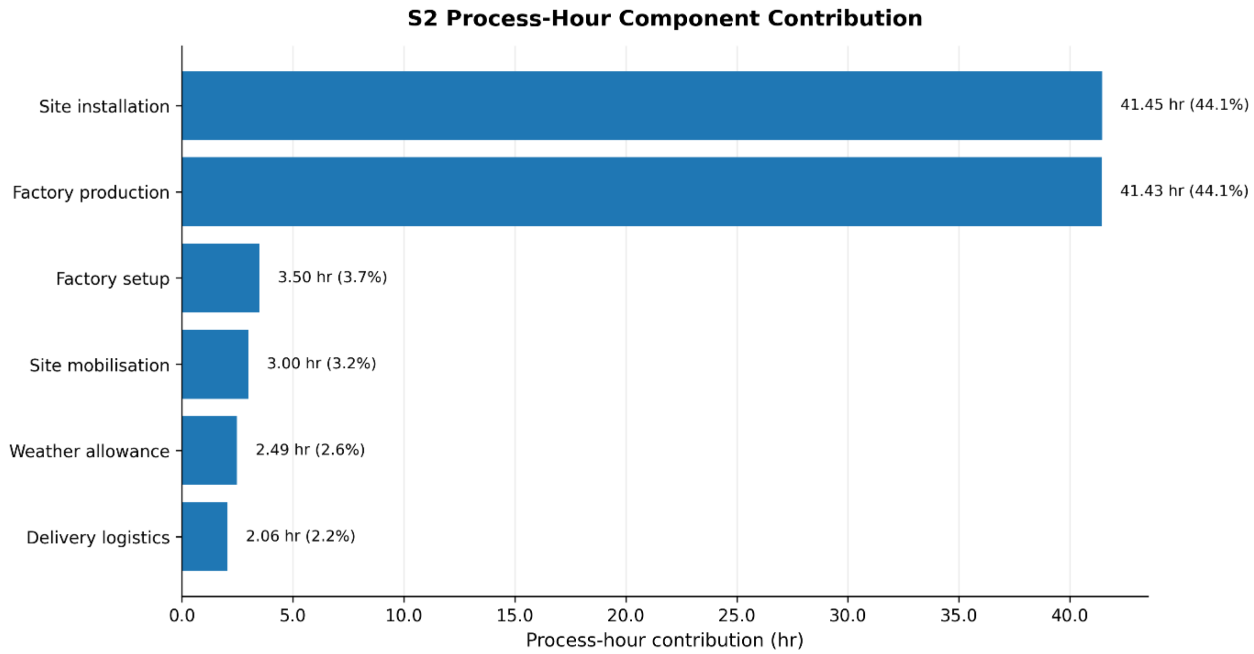

Figure 4-4 S2 process-hour component contribution from the implemented CFS framing time model. The figure shows the factory, logistics, site, and weather-related components that form the model-calculated S2 process-hours.

Table 4-8 Objective-function implementation traceability.

Objective	Implemented function	Component output	Research verification role
CFS framing process cost	calculate_cost_components_abs ()	Cost component breakdown output	Verifies that total cost is assembled from material, factory, setup, site labour, logistics, opening/detailing, waste, and overhead components
CFS framing A1–A5 carbon	calculate_carbon_components_abs ()	Carbon component breakdown output	Verifies that total carbon is assembled from A1–A3 steel, waste/offcut carbon, A4 transport, factory energy, and A5 installation components
CFS framing process-hours	calculate_time_components_abs ()	Time component breakdown output	Verifies that total time is assembled from factory production, setup, delivery logistics, site mobilisation, site installation, and weather allowance
Component export	build_component_breakdown_tables ()	Cost, carbon, and time breakdown outputs	Provides reproducible component-level evidence for S2 and representative solution interpretation

The implementation traceability of the three objective functions is verified through the component-breakdown outputs generated during the reported model run. The cost, carbon, and time objectives are not exported only as total values; each is decomposed into its governing process components. This confirms that the mathematical formulation developed in Chapter 3 was implemented as a component-level engineering response model and that another researcher can inspect the reported outputs to verify how each objective value was assembled.

4.7: Material Weight and Interaction Effects

Material weight and interaction diagnostics were included to check whether the implemented response model behaves consistently with the engineering assumptions defined in Chapter 3. These diagnostics do not introduce new optimisation objectives. Instead, they inspect how rationalisation, opening/detailing complexity, and prefabrication level affect steel mass, cost, carbon, and process-hour response terms within the fixed case-study boundary.

Steel mass is treated as a bounded response rather than as a freely optimised structural variable. The fixed case-study geometry, wall-panel quantity, roof-truss quantity, CFS section, steel thickness, grade, and coating remain unchanged during optimisation. As a result, the model does not redesign structural member sizes or claim structural capacity optimisation. Instead, the material-weight response represents smaller production-related effects: offcut and waste reduction from repetition, modest prefabrication precision benefits, grouping or oversizing penalties from excessive rationalisation, and opening/detailing-related steel effects.

Panel rationalisation intensity R_p and truss rationalisation intensity R_t influence the model through competing repetition and grouping mechanisms. Moderate rationalisation can reduce the number of effective panel or truss families, improve repetition, and reduce setup, waste, and site-handling effort. However, aggressive rationalisation may force larger or less efficient grouped elements, increasing oversizing, handling, logistics, or fabrication penalties. This is important because the model does not assume that maximum rationalisation is automatically optimal. Instead, the preferred level of rationalisation emerges from the balance between repetition benefits and over-standardisation penalties.

Opening/detailing complexity OC is implemented as a design-for-manufacture and detailing response, not as a change in the number of architectural openings. Lower OC represents more standardised lintels, jamb studs, trimming details, service penetrations, and repeated opening modules, while higher OC represents custom detailing, additional trimming, local reinforcement, coordination effort, and installation difficulty. These effects influence material weight, cost, carbon, and time because opening details affect not only steel usage but also factory preparation and site installation effort.

Prefabrication level P_f affects the model through both beneficial and adverse mechanisms. Higher prefabrication can improve cutting precision, reduce site installation effort, reduce waste, and lower site-stage installation carbon. At the same time, higher P_f can increase factory processing, factory production time, delivery coordination, transport sensitivity, and overload effects when prefabrication approaches the upper end of the tested range. This factory–site trade-off is central to the model because prefabrication is not treated as a simple one-directional saving.

The interaction terms are therefore used to represent engineering trade-offs rather than arbitrary numerical penalties. The nondominated front is shaped by the combined effect of material-weight response, repetition benefits, oversizing penalties, opening/detailing complexity, prefabrication response, logistics sensitivity, and factory overload. This structure allows later results to be interpreted as construction and fabrication trade-offs: a solution may improve site time but increase factory cost, reduce carbon through lower A5 installation effort but increase factory-related effects, or reduce setup through rationalisation while introducing grouping penalties. The interaction model is therefore essential for translating the optimisation output into civil and structural engineering decision support. The implementation traceability for this section is provided by the material-weight response, engineering-state calculation, opening-complexity interaction plot, and cost–time prefabrication breakeven plot. Together, these outputs verify that the objective trade-offs arise from the coupled response structure of the model and not from independent objective scaling.

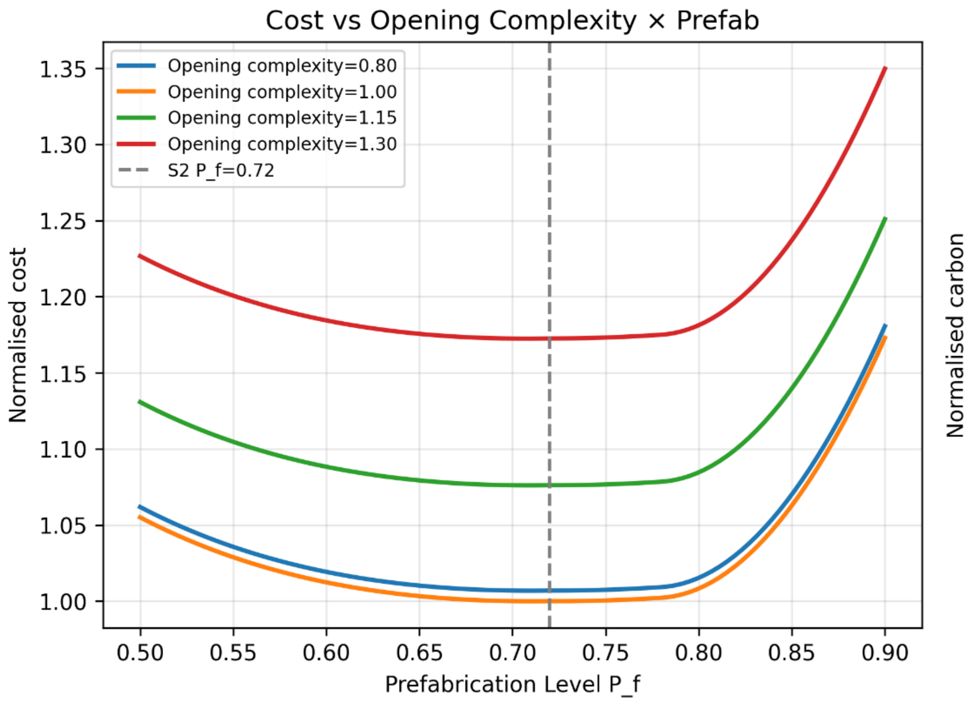

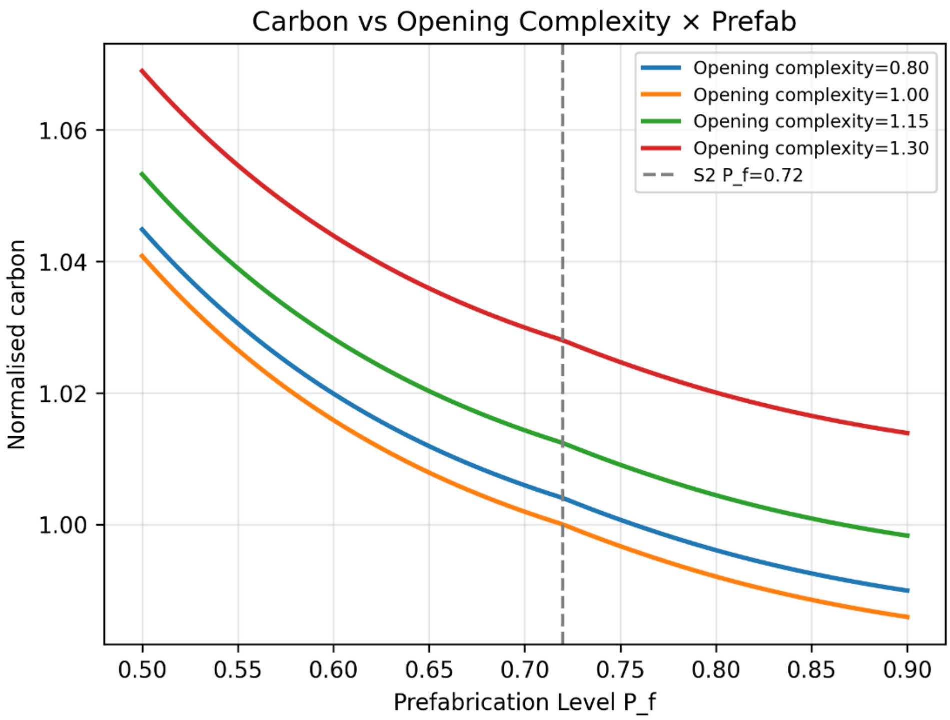

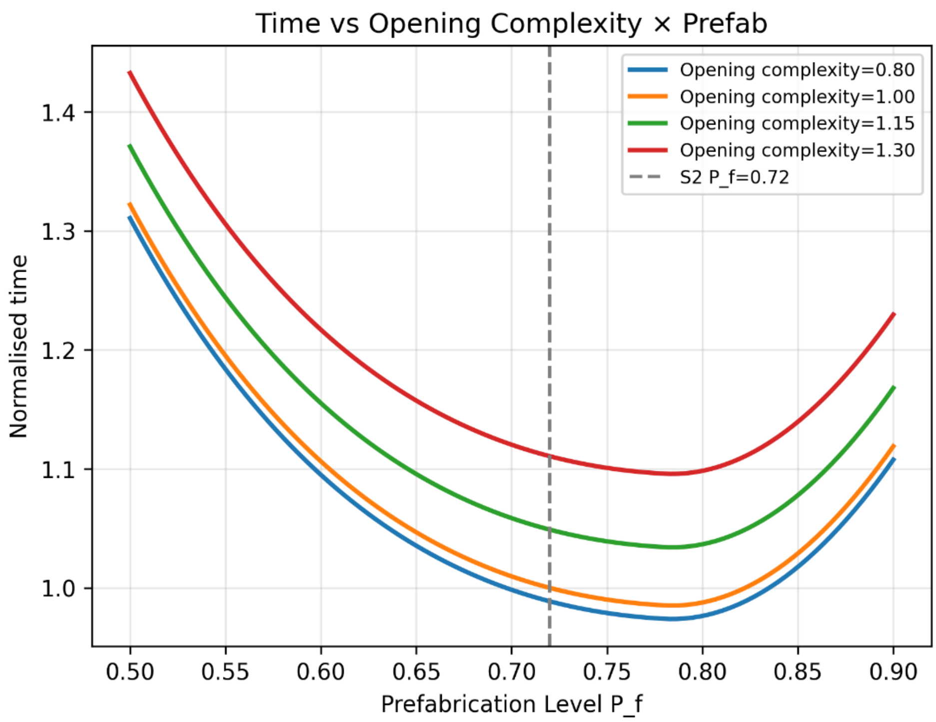

Figure 4-5 Opening/detailing complexity and prefabrication interaction response under the S2 rationalisation reference condition.

Figure 4-5 shows the interaction between opening/detailing complexity and prefabrication level under the S2 rationalisation condition. The three subplots demonstrate that the effect of prefabrication is not constant across all opening/detailing conditions. For cost, lower opening complexity produces the most favourable response, while higher opening complexity shifts the cost to curve upward and increases the penalty at higher prefabrication levels. This indicates that prefabrication benefits depend on the extent to which opening details, lintels, jamb studs, trimming, and service penetrations are standardised before fabrication.

These diagnostic plots should be interpreted as response-model checks rather than independent empirical validation. They confirm that the implemented equations produce the intended engineering behaviour, including diminishing prefabrication benefits, nonlinear opening/detailing penalties, and bounded rationalisation effects.

The carbon and time responses show the same interaction logic, but with different sensitivities. Carbon generally decreases with increasing prefabrication because the model allows site-stage carbon and waste effects to be reduced as more work is transferred to controlled factory conditions. However, higher opening complexity still maintains a higher carbon response because additional trimming, local reinforcement, coordination, and installation effort remain embedded in the framing process. The time response also shows that prefabrication can reduce process-hours up to a practical range, but high opening complexity weakens this benefit and raises the time curve. This confirms that the implemented model treats opening/detailing complexity as a coupled fabrication and installation variable rather than as an isolated penalty.

Overall, Figure 4-5 supports the modelling assumption that the Pareto front is shaped by interaction effects between prefabrication and detailing complexity. The figure demonstrates that high prefabrication alone does not guarantee the best cost, carbon, or time outcome; its effectiveness depends on whether the framing configuration is also rationalised for manufacture and installation. This is important for CFS prefabrication practice because design-for-manufacture decisions made during detailing can materially affect factory productivity, site installation effort, and the resulting cost–carbon–time balance.

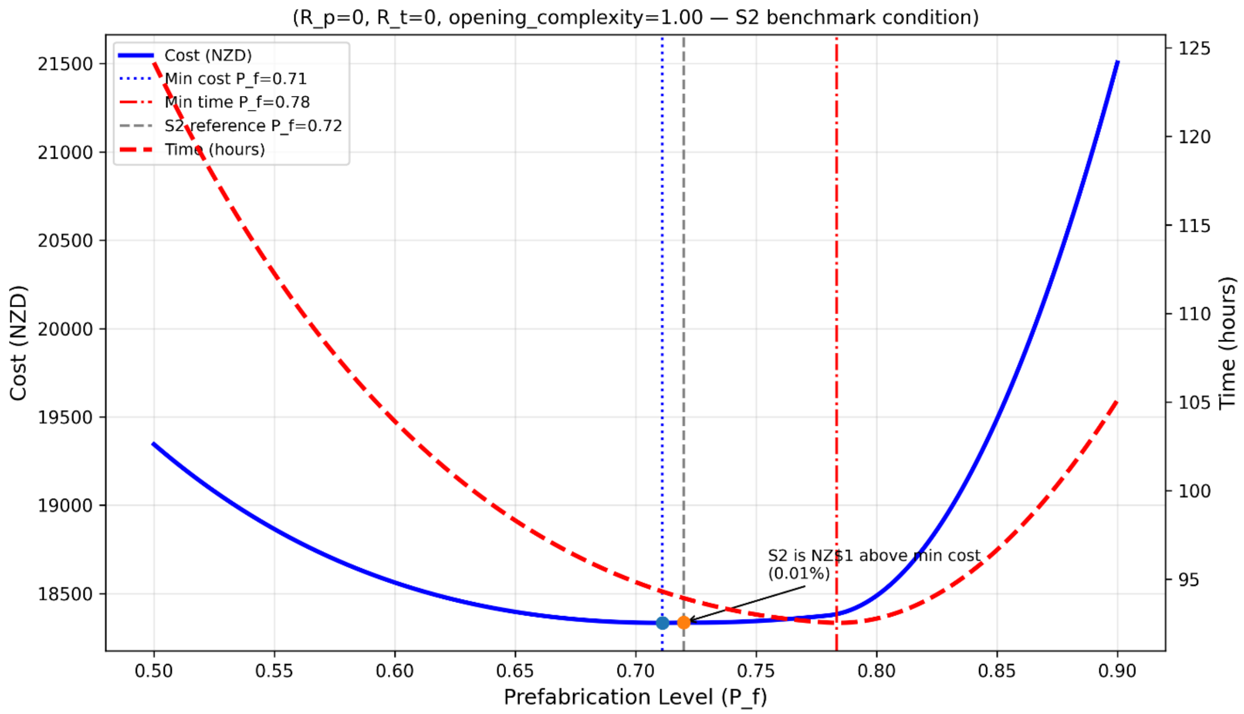

Figure 4-6 Cost–process-hour response across prefabrication level under the S2 rationalisation condition.

Figure 4-6 illustrates the relationship between cost and process-hours as prefabrication level changes while panel rationalisation, truss rationalisation, and opening complexity are held at the S2 reference condition. The cost curve shows a shallow minimum around the medium-prefabrication range, while the time curve continues to improve until a slightly higher prefabrication level before increasing again. This difference between the minimum-cost and minimum points demonstrates that prefabrication is not implemented as a simple one-directional saving. Instead, it acts as a factory–site trade-off.

The figure shows that increasing prefabrication can reduce site-related process-hours because more work is shifted from site installation to factory production. However, beyond the favourable range, higher prefabrication increases cost and eventually time due to additional factory processing, logistics coordination, delivery complexity, and overload effects. The S2 reference point sits close to the cost-optimal region, while the minimum point occurs at a higher prefabrication level. This indicates that a decision-maker seeking lower process-hours may need to accept a different cost position than one seeking minimum cost.

Figure 4-6 therefore provides implementation evidence that the model captures a realistic prefabrication trade-off. The response is not produced by arbitrary penalties, but by the balance between factory production effort, site installation reduction, logistics burden, and prefabrication-related process effects. This supports the later Pareto-front interpretation because the optimiser is not simply selecting the highest prefabrication level; it is searching for configurations where the factory, logistics, site labour, and time effects remain balanced within the fixed CFS framing boundary.

Together, Figures 4-5 and 4-6 verify that the implemented objective functions are governed by coupled engineering response mechanisms. Opening/detailing complexity changes the effectiveness of prefabrication, while prefabrication itself produces competing cost and process-hour effects. These results support the interpretation that the final nondominated front arises from material-weight response, repetition benefits, grouping penalties, detailing complexity, factory–site transfer, and logistics interaction, rather than from arbitrary objective scaling.

The engineering assumptions used in the response model are intended to be transparent rather than universal. Parameters directly tied to market rates and emission factors are assigned uncertainty ranges and propagated through Monte Carlo analysis. Parameters defining the fixed S2 process structure, such as base factory panel and truss production time, are treated as case-study reference coefficients and reported with plausibility ranges in the calibration audit. Low-impact carbon coefficients, such as factory average power, are disclosed and bounded because their contribution to total S2 carbon is small. This distinction allows the model to remain reproducible and bounded without overstating the level of empirical validation available for every coefficient.

4.8: Model Verification and S2 Reasonableness Evidence

The validation and diagnostic stage checks whether the implemented framework produces internally consistent, feasible, and interpretable outputs before the optimisation results are discussed. The purpose of these checks is not to provide external commercial validation, but to verify implementation consistency, feasibility screening, S2-reference reproduction, objective normalisation, Pareto extraction, and S2-relative improvement classification.

The reported model run used Latin Hypercube Sampling initialisation, a population size of 200, 300 generations, one thousand Monte Carlo simulations, and the multi-seed set 11,22,33,44,55. The final Pareto-front output contained 8302 nondominated solutions. All reported Pareto-front solutions were recorded as feasible within the implemented CFS framing constructability boundary, with no hard or soft constraint violations reported for the interpreted solution set. This confirms that the reported alternatives remain within the defined case-study scope rather than relying on infeasible or out-of-bound configurations.

A key validation requirement was the distinction between Pareto membership and genuine improvement relative to S2. A solution was not treated as superior to S2 simply because it appeared on the nondominated front. Instead, the representative solutions were checked objective by objective against the S2 reference for CFS framing process cost, CFS framing A1–A5 carbon, and CFS framing process-hours. This prevents the research from overstating trade-off solutions as universal improvements.

The All-objective improvement and trade-off check showed that the cost-minimum solution reduced cost by 1.61%, carbon by 1.08%, and process-hours by 3.70% relative to S2. The time-minimum solution also improved all three objectives, reducing cost by 0.50%, carbon by 1.35%, and process-hours by 6.76%. In contrast, the carbon-minimum solution reduced carbon by 2.59% but increased cost by 15.58% and process-hours by 8.25%. The knee/balanced solution reduced carbon by 2.12% and process-hours by 4.23% but accepted a 0.60% cost premium. These classifications are important because they separate dominant alternatives from specialist low-carbon and balanced trade-off alternatives.

An additional diagnostic check was conducted by fixing the opening/detailing complexity index at the S2 condition, $OC = 1.0$. This test is important because opening/detailing complexity represents the standardisation of lintels, jamb studs, service penetrations, and trimming details, and therefore has a strong design-for-manufacture influence on CFS framing performance. With OCfixed at the S2 value, 2,695 of the 8302 Pareto-front solutions still dominated S2 across cost, carbon, and process-hours, while 5608 did not. This result shows that a meaningful subset of improvement remains when opening/detailing complexity is controlled, but it also confirms that opening/detailing standardisation is a major contributor to the reported trade-off structure.

The variable-activity diagnostics were used to check whether the design variables remained meaningful across the Pareto front. Panel rationalisation intensity R_p, truss rationalisation intensity R_t, panel scheme, opening/detailing complexity OC, and prefabrication level $P_f$all showed active variation across the nondominated solution set. The truss scheme, however, collapsed to a single selected scheme. This is not interpreted as a modelling failure. Within the fixed case-study boundary, it indicates that one truss standardisation branch was consistently preferred after balancing repetition benefits, span-family grouping effects, material response, logistics, and process-hour implications.

Overall, the diagnostic stage strengthens the engineering credibility of the framework. The reported Pareto front is not treated as a raw optimisation result; it is interpreted as a checked decision set for CFS framing rationalisation, prefabrication, and detailing decisions. This supports the research aim by showing that the decision-support framework can distinguish between dominant solutions, specialist low-carbon alternatives, and balanced trade-off configurations while remaining transparent about modelling boundaries and limitations.

The reported optimisation run used the same feasibility-screening thresholds defined in Section 3.11. Hard feasibility screening was applied to the practical CFS framing limits of maximum panel width ($b_{panel} \leq 2.40, \mathrm{m}$), maximum panel height ($h_{panel} \leq 3.00, \mathrm{m}$), maximum stud spacing ($s_{stud} \leq 600, \mathrm{mm}$), minimum steel thickness ($t_{steel} \geq 0.55, \mathrm{mm}$), and maximum estimated lift mass ($m_{lift} \leq 2000, \mathrm{kg}\}$). The fixed case-study section uses 0.75 mm steel thickness, so the thickness condition was satisfied by the adopted structural boundary. Soft constructability screening was also retained for preferred unique truss spans ($N_{t,unique} \leq 25$), minimum factory repetition score ($R_{factory} \geq 0.65$), and preferred opening/detailing complexity ($OC \leq 1.20$). Hard-threshold violations were treated as infeasible through the objective-space penalty structure, whereas the soft constructability conditions were retained as preference-based penalties rather than absolute rejection rules. All final Pareto-front solutions reported in Chapter 5 satisfied the hard feasibility thresholds within the implemented CFS framing constructability boundary.

Table 4.9 summarises the validation and diagnostic evidence used to support the reported model run. The checks include S2 reference reproduction, feasible-solution counts, constraint status, grid-reference comparison, hypervolume convergence, multi-seed stability, component-breakdown consistency, and dashboard data export. Together, these checks establish internal model traceability and computational reproducibility.

Table 4-9 Verification, diagnostic, and practical reasonableness evidence for the reported model run.

Validation item	**Reported output evidence**	**Reported finding**	**Civil/structural engineering purpose**
Run traceability	run_manifest.csv	LHS initialisation; population = 200; generations = 300; Monte Carlo runs = 1000; multi-seed set = (11, 22, 33, 44, 55)	Confirms the computational basis of the reported research run and separates it from earlier test runs
Pareto-front size	pareto_front.csv	8302 nondominated solutions	Shows that interpretation is based on a substantial solution set, not a single isolated configuration
Feasibility status	pareto_front.csv and representative_solutions.csv	Reported Pareto and representative solutions are feasible within the implemented CFS framing boundary	Confirms that interpreted alternatives remain inside the constructability limits used by the model
S2 benchmark status	validation_status.csv	S2 is a model-calculated reference, not a measured commercial benchmark	Prevents the research from presenting modelled S2 values as contractor tender or measured project performance
Time boundary	validation_status.csv	Time is reported as CFS framing process-hours only	Prevents process-hour results from being misread as full construction programme duration
Carbon boundary	validation_status.csv	Carbon is limited to CFS framing A1–A5 only	Prevents the carbon results from being misread as whole-building life-cycle carbon
Fixed structural scope	validation_status.csv	Geometry, framing quantities, steel section, member sizing, and structural verification are fixed or outside scope	Confirms that the framework is a configuration decision-support model, not a structural design optimiser
Honest S2 improvement check	honest_improvements_vs_s2.csv	Cost-min and time-min improve all three objectives; carbon-min improves carbon only; knee/balanced improves carbon and time with a small cost premium	Separates dominant alternatives from specialist and trade-off solutions
Opening-complexity fixed check	oc_fixed_comparison.csv	2,695 solutions still dominate S2 when (OC=1.0); 5608 do not	Tests whether improvement depends on reduced opening/detailing complexity
Variable-activity check	pareto_degeneracy_report.csv	(R_p), (R_t), panel scheme, (OC), and (P_f) remain active; truss scheme collapses	Identifies which CFS framing decisions remain trade-off active and which converge to a preferred case-study setting
Design-variable interpretation	variable_dominance_finding.csv	Rationalisation, opening/detailing complexity, and prefabrication level generate trade-off behaviour	Converts numerical variable spread into practical DfMA interpretation for CFS framing decisions

Validation item	Reported output evidence	Reported finding	Civil/structural engineering purpose
Representative solution audit	representative_solutions.csv	S2, cost-min, carbon-min, time-min, and knee/balanced alternatives retained	Provides a manageable decision set for later Pareto-front, stakeholder-ranking, and dashboard discussion

Table 4-9 demonstrates that the reported optimisation outputs were checked from both numerical and engineering perspectives. The run traceability and Pareto-front size confirm the basis of the reported solution set. The feasibility and scope checks confirm that the reported alternatives remain within the fixed CFS framing boundary. The honest S2 comparison prevents the research from overstating trade-off solutions as universally superior alternatives. The OC-fixed comparison and variable-activity diagnostics further show whether the front is shaped by genuine rationalisation, prefabrication, and detailing interactions or by dependence on a narrow modelling condition.

This validation structure is suitable for the research because it aligns directly with the research aim. The validation checks confirm that the implemented outputs remain consistent with the scope and numerical definitions established in Chapter 3.it is presented as a transparent CFS framing decision-support model that helps designers and fabricators compare configuration alternatives under fixed case-study assumptions. The diagnostic outputs therefore provide the evidence needed to interpret the reported Pareto front responsibly.

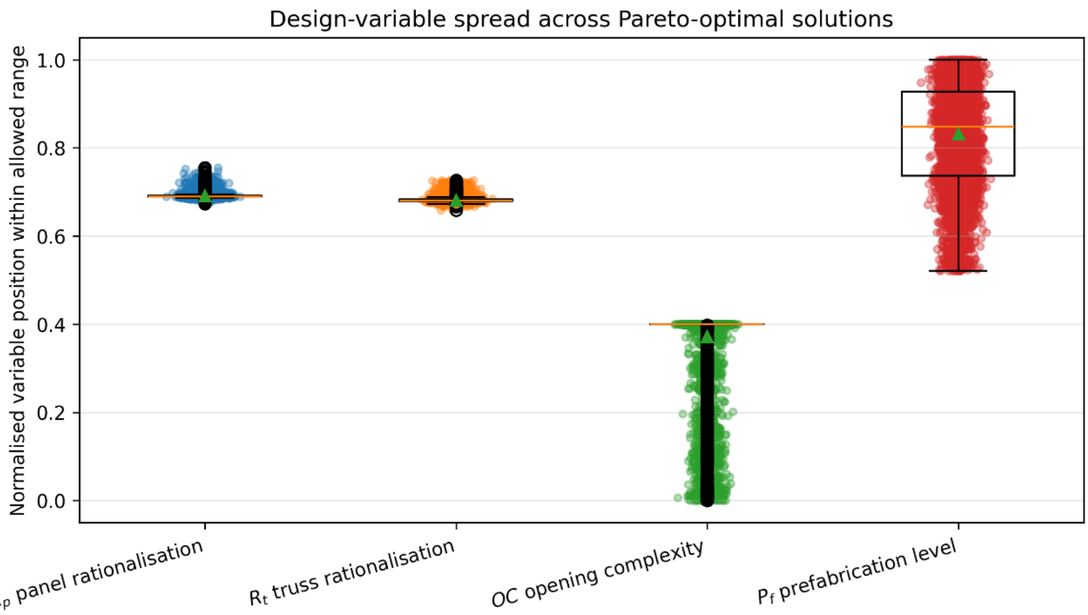

Figure 4-7 Pareto-front variable spread for the reported optimisation run.

Figure 4-7 shows the spread of the design variables across the nondominated CFS framing solutions. The values are normalised within their allowable ranges, so the figure indicates the relative position of each variable within the tested design space. This diagnostic is used to identify which configuration decisions remain active in the cost–carbon–time trade-off and which variables show convergence toward a preferred case-study range.

Panel rationalisation $R_p$and truss rationalisation $R_t$are concentrated within narrow upper-range bands. From a CFS fabrication perspective, this indicates a stable preference for moderate-to-high rationalisation, where repetition benefits can be achieved without pushing the system into excessive grouping or handling penalties. This is a useful engineering finding because it suggests that the Pareto-optimal configurations favour practical standardisation rather than either highly bespoke framing or extreme rationalisation.

Opening/detailing complexity $OC$is concentrated toward the lower part of the allowable range. This indicates that the nondominated solutions favour more standardised opening and trimming details. In practical CFS framing terms, repeated lintel arrangements, jamb-stud details, trimming layouts, and service-penetration treatments improve constructability by reducing detailing variation, fabrication coordination, and site installation effort. The result therefore supports the design-for-manufacture interpretation that opening/detailing standardisation is a key contributor to improved framing performance.

Prefabrication level $P_f$shows the widest spread across the nondominated front. This confirms that prefabrication remains a genuine trade-off variable rather than converging to a single value. Higher prefabrication can reduce site installation effort and improve production control, but it can also increase factory processing, transport coordination, and logistics burden. The broad $P_f$range therefore reflects the practical factory–site balance that CFS fabricators and installers must manage.

Overall, Figure 4-7 interpreted as evidence of meaningful design-variable behaviour, not as an optimisation weakness. The Pareto front is shaped by a combination of stable rationalisation preferences, strong benefits from opening/detailing standardisation, and a broad prefabrication trade-off range. This supports the research aim by showing that the framework produces interpretable CFS framing configuration guidance rather than a purely abstract optimisation output.

4.8.1: External Plausibility Check for Model-Calculated S2 Values

The external plausibility check of the S2 reference values is promoted here because it is the main evidence that the implemented coefficient set produces reasonable absolute outputs before optimisation. The S2 values are model-calculated rather than measured project records, but they must still be plausible relative to the adopted steel mass, cost rates, carbon factors, factory and site process assumptions, transport allowance, and A5 construction-stage assumptions. Table 4-10 consolidates the evidence basis for the three S2 objective values used throughout the study.

Table 4-10 External plausibility check for model-calculated S2 reference values.

S2 output	Model-calculated value	Plausibility reference or benchmark basis	Interpretation
CFS framing process cost	NZ$18,335	Compare against NZ CFS framing cost rates, steel supply/fabrication rates, site labour rate assumptions, setup/logistics allowances, and S2 steel mass of 2,195.264 kg	Used as a case-study process-cost benchmark, not a contractor tender price

S2 output	Model-calculated value	Plausibility reference or benchmark basis	Interpretation
CFS framing A1–A5 carbon	5,367 $kgCO_2e$	Compare against steel mass × steel emission factor, A4 transport factor, A5 installation allowance, and waste/offcut assumptions	Plausible within A1–A5 CFS framing boundary; not certified whole-building LCA
CFS framing process-hours	93.9 h	Compare against factory panel/truss preparation, delivery/logistics, mobilisation, and site installation process-hour assumptions	Plausible as process effort; not full construction programme duration

Table 4-10 does not validate the model externally. Instead, it shows that the S2 cost, carbon, and process-hour values are plausible within the adopted case-study boundary and coefficient set. This distinction is important because factory labour records, site productivity records, contractor tender prices, and project-specific verified carbon data were not available. The S2 check therefore supports reasonableness and traceability, while the optimisation results in Chapter 5 and uncertainty results in Chapter 6 remain conditional on the implemented response model.

The external reasonableness check was used as a practical plausibility screen for the model-calculated S2 objective magnitudes. This check does not convert the model into a supplier quotation, certified embodied-carbon assessment, or measured productivity record. Instead, it tests whether the S2 cost, carbon, and process-hour values are reasonable for the narrower CFS framing process boundary used in this study.

For cost, the model-calculated S2 CFS framing process cost is NZD 18,335.47. Normalised by the case-study envelope/framing area of 217.2555 m², this corresponds to 84.39 NZD/m² of CFS framing envelope area. Normalised by the footprint area of 98.292 m², it corresponds to 186.54 NZD/m² of building footprint. These values are interpreted as CFS framing process-cost indicators only, because the model excludes foundations, linings, cladding, services, consent costs, fit-out, contractor margin, and other whole-house cost components. The cost magnitude is therefore practically plausible as a bounded CFS framing process-cost reference, but it should not be interpreted as a complete market price for a house or as a supplier tender. No independent quantity-surveyor estimate, or completed-project CFS supplier quotation was available for this exact framing scope; therefore, the cost check is treated as dimensional and order-of-magnitude plausibility rather than external price validation.

For carbon, the model-calculated S2 CFS framing A1–A5 carbon is 5,366.79 $kgCO_2e$. Relative to the envelope/framing area, this corresponds to approximately 24.70 $kgCO_2e/m^2$ of CFS framing envelope area. Relative to the footprint area, it corresponds to approximately 54.60 $kgCO_2e/m^2$ of building footprint. The dominant component is A1–A3 steel carbon, which is governed by the S2 steel mass of 2,195.264 kg and the adopted steel emission factor. The carbon magnitude is therefore physically consistent with the steel-mass-driven A1–A3 boundary, with smaller contributions from A4 transport, factory energy, A5 installation, and waste allowance. This supports the reasonableness of the S2 carbon value for comparative A1–

A5 CFS framing assessment, while recognising that a project-specific EPD and full building LCA would be required for certified carbon reporting.

For process-hours, the model-calculated S2 value is 93.92 h. This corresponds to approximately 0.43 h/m² of envelope/framing area, 0.96 h/m² of footprint area, or 0.89 h per framing element when distributed across the 48 wall panels and 58 roof trusses. This value is interpreted as a process-hour estimate for CFS framing production, setup, delivery/logistics, mobilisation, site installation, and weather allowance. It is not a full construction programme and does not include procurement lead time, inspection delays, non-CFS trades, subcontractor sequencing, or calendar float. The value is therefore plausible as a bounded process-effort benchmark, but it should not be read as total project duration.

Overall, the S2 reasonableness check supports the use of S2 as a consistent model-calculated reference for normalisation, Pareto comparison, Monte Carlo paired testing, and dashboard interpretation. The check strengthens practical credibility by showing that the S2 absolute values are dimensionally traceable and compatible with the stated CFS framing boundary. However, it remains a plausibility check rather than full empirical validation. Stronger external validation would require measured fabrication hours, site installation records, supplier pricing, transport records, and project-specific carbon declarations from completed CFS projects. The absence of measured factory labour, site labour, supplier pricing, and project-specific carbon declaration data remains the main external-validation limitation of the study.

The S2 reasonableness check was also compared qualitatively against the external benchmark sources reviewed earlier in the thesis. The cost value of NZD 18,335.47 is not compared with whole-house construction rates from QV CostBuilder on a like-for-like basis because the model excludes foundations, cladding, linings, services, fit-out, consent costs, contractor margin, and non-CFS trades. Instead, the comparison indicates that the S2 result is plausible only as a CFS framing process-cost magnitude. Similarly, the S2 carbon value of 5,366.79 kgCO_2e is consistent with the steel-mass-driven A1–A3 logic and the CFS framing A1–A5 carbon benchmark structure reviewed in Table 2-5, but it is not a product-specific EPD or certified building LCA value. The S2 process-hour value of 93.92 h is also plausible as a bounded process-effort estimate because it is assembled from factory production, setup, logistics, mobilisation, site installation, and weather allowance rather than from a full CPM programme. Therefore, the external check supports practical plausibility within the stated CFS framing boundary, but it does not constitute full empirical validation against measured commercial project records.

4.9: Dashboard Export System

The dashboard export system converts the final model outputs into a stable data structure for Chapter 7 interpretation. The export stage does not generate new optimisation results. Its role is to transfer the reported Pareto front, representative solutions, S2 reference values, sensitivity outputs, Monte Carlo robustness results, TOPSIS rankings, and model-audit information into dashboard-ready files.

The dashboard export system is intentionally one-way for the reported study results. The dashboard reads the exported CSV, JSON, and figure files generated by the optimisation and post-processing stages; it does not rerun NSGA-II, regenerate the official Pareto front, or replace the reported optimisation outputs. This boundary is important because the dashboard is used as an interpretation and decision-support interface, not as a live optimisation engine or certified engineering software. Scenario controls in the dashboard should therefore be interpreted as bounded stress-testing and communication tools rather than as replacement optimisation runs.

The dashboard export includes the Pareto-front solutions, representative configurations, S2 comparison data, sensitivity results, Monte Carlo robustness outputs, TOPSIS ranking, and key scope labels. These outputs allow the dashboard to support practical review of alternative CFS framing configurations under different stakeholder priorities. The dashboard therefore functions as an interpretation layer for the optimisation results rather than as a structural design tool, contractor pricing system, or whole-building life-cycle assessment platform.

For reproducibility, the dashboard dataset is written to the dashboard/data folder. The live dashboard file is stored as dashboard_config. json. This ensures that the dashboard used for research interpretation can be traced back to the same optimisation run, S2 reference, objective definitions, and uncertainty assumptions used in the reported results.

Table 4-11 Dashboard export mapping.

Export item	Output evidence	Dashboard role	Engineering interpretation
Dashboard configuration	dashboard_config. json	Main dashboard data source	Preserves the reported optimisation dataset for review
Pareto-front solutions	pareto_front.csv	Pareto explorer	Allows comparison of cost, carbon, and process-hour trade-offs
Representative solutions	representative_solutions.csv	Solution cards	Summarises cost-minimum, carbon-minimum, time-minimum, and balanced alternatives
S2 comparison	S2 improvement outputs	S2 comparison panel	Shows whether alternatives improve or trade off against the reference case
Sensitivity results	Sensitivity output tables	Design-driver panel	Identifies which variables most affect the framing response
Monte Carlo robustness	Monte Carlo robustness outputs	Uncertainty panel	Shows whether selected solutions remain dependable under uncertainty
TOPSIS ranking	TOPSIS ranking outputs	Stakeholder-ranking panel	Supports different priorities between cost, carbon, and time
Imported Pareto front / scenario controls	Pareto and scenario output	Used for filtering, comparison, and bounded scenario interpretation	Dashboard does not rerun NSGA-II or replace the official imported Pareto front
Scope labels	Validation-status outputs	Boundary notes	Prevents misinterpretation as structural verification, tender pricing, or whole-building LCA

The dashboard data export preserves the distinction between reported optimisation outputs and later dashboard interaction. The imported Pareto front remains the final optimisation result generated by the model run. Dashboard filtering, ranking, and scenario controls support interpretation and bounded stress-testing, but they do not replace the reported NSGA-II run or create a new certified optimisation result.

4.10: Reproducibility and File Management

Reproducibility was supported by recording the run configuration, random seeds, output folders, generated tables, figure files, and dashboard-ready data bundle for the reported model run. The purpose of this section is to identify the records required to trace the numerical results reported in Chapters 5–7 back to the final implemented workflow.

The output structure separates tabulated evidence, figures, and dashboard data. Research tables are retained in results_fixed/tables, research figures are retained in results_fixed/figures, and dashboard-ready files are retained in dashboard/data. This folder separation supports transparent review because the source of each reported table, figure, and dashboard panel can be traced back to the same run timestamp. The reported-run folder should remain unchanged after final research figures and tables are produced. Any later rerun is treated as a new model execution with a new timestamp rather than replacing the reported evidence.

The run settings are documented through the run manifest and associated output files. For the reported run, the framework used LHS initialisation, a population size of 200, 300 generations, 1000 Monte Carlo simulations, and the multi-seed set 11,22,33,44,55. These settings are part of the reproducibility record because stochastic optimisation and uncertainty analysis can produce different outputs if seed values, population size, or run settings are changed.

Figure and table management was also standardised for research use. CSV outputs were retained as the primary numerical record, while figures were exported as high-resolution image files suitable for insertion into the research document. The final figure set was inserted directly from the archived results_fixed/figures folder so that the visual evidence remains consistent with the numerical outputs in results_fixed/tables. This avoids mismatch between reported values, plots, and dashboard interpretation.

The computational environment was recorded with the reported-run archive. At minimum, the environment record should include the Python version and the versions of NumPy, pandas, matplotlib, DEAP, SciPy, and scikit-learn. This is necessary because optimisation behaviour, numerical processing, and plot formatting can change between package versions. Recording the environment allows another researcher to rerun the framework under comparable computational conditions and verify the same CFS framing decision-support outputs.

Reproducibility in this chapter refers to reproducing the reported model run and its exported evidence under the stated case-study assumptions. It does not imply that the model has been externally validated against independent contractor quotations, measured site productivity records, product-specific EPDs, or multi-project CFS datasets. These external validation tasks are outside the scope of the study and are identified as future research requirements.

Any later change to the model files, dashboard files, or repository should be treated as a new version unless a new run manifest and output archive are generated. The numerical results reported in this study remain tied to the documented final model run and its exported tables, figures, and dashboard data files.

Chapter 5: NSGA-II Optimisation Results

Chapter 5 presents the optimisation results from the reported LHS-initialised NSGA-II model run. The purpose is to evaluate whether the implemented response model generated a stable, feasible, and interpretable nondominated archive for early-stage CFS framing cost–carbon–process-hour screening. The chapter first reports convergence, stability, LHS initial-population coverage, solution-count auditing, and grid-reference comparison. It then interprets the final archive through S2-relative classification, representative solutions, derived panel and truss reporting categories, component changes, and TOPSIS-based stakeholder ranking.

The chapter distinguishes carefully between Pareto membership and S2-relative improvement. A solution can belong to the nondominated archive without improving all objectives relative to S2. Conversely, a subset of nondominated solutions may improve CFS framing process cost, A1–A5 carbon, and process-hours simultaneously if the S2 reference is not locally efficient under the implemented response model. This distinction is central to the interpretation of the 8,302 nondominated solutions, including the 2,695 all-objective improvement solutions and the 5,607 trade-off solutions.

5.1: Convergence and Stability

The optimisation output was first assessed through convergence, stability, feasibility-screening, and solution-count diagnostics before the Pareto-front results were interpreted. These diagnostics do not prove recovery of the exact global Pareto set; rather, they support the use of the reported nondominated archive as an internally stable NSGA-II approximation under the stated model settings. The diagnostic assessment used hypervolume convergence, grid-reference IGD, multi-seed repeatability, and a solution-count audit to confirm that the final nondominated archive was suitable for subsequent CFS framing decision-support analysis.

The hypervolume history for the reported run is summarised in Table 5-1 and plotted in Figure 5-1. The hypervolume trend indicates that the reported run reached a stable nondominated approximation early and then improved only marginally in later generations. The final hypervolume was $HV = 0.04233043$, and the recorded convergence pattern indicates that most of the final objective-space coverage was achieved well before the final generation. This supports the interpretation that the final Pareto front was not produced by late-generation stochastic fluctuation.

Table 5-1 Hypervolume progress for the reported NSGA-II run.

Generation	Hypervolume	HV as % of final HV	Nondominated archive size	Physical archive size	HV change
10	0.04177721	98.69	89	84	0.00000000
20	0.04216720	99.61	386	349	0.00038999
30	0.04225143	99.81	627	562	0.00008424
50	0.04231879	99.97	1,208	1,031	0.00004567
100	0.04232900	100.00	2,695	1,983	0.00000018
150	0.04232998	100.00	4,242	2,779	0.00000009
170	0.04233007	100.00	4,793	3,000	0.00000009
200	0.04233034	100.00	5,642	3,000	0.00000018
210	0.04233043	100.00	5,921	3,000	0.00000009

Generation	Hypervolume	HV as % of final HV	Nondominated archive size	Physical archive size	HV change
220	0.04233043	100.00	6,198	3,000	0.00000000
250	0.04233043	100.00	6,971	3,000	0.00000000
300	0.04233043	100.00	8,302	3,000	0.00000000

The convergence history also indicates that the 300-generation setting was conservative for the implemented response model. By generation 100, the hypervolume had already reached the reported final value to two decimal places, and from generation 210 onward no further HV change was observed at the reported precision. The later generations therefore mainly increased the density of nondominated alternatives within the same trade-off region rather than discovering a materially different objective-space boundary. The 300-generation setting was retained to provide a stable archive for representative-solution selection, robustness testing, and dashboard filtering, not because substantial objective-space improvement continued to occur after generation 100.

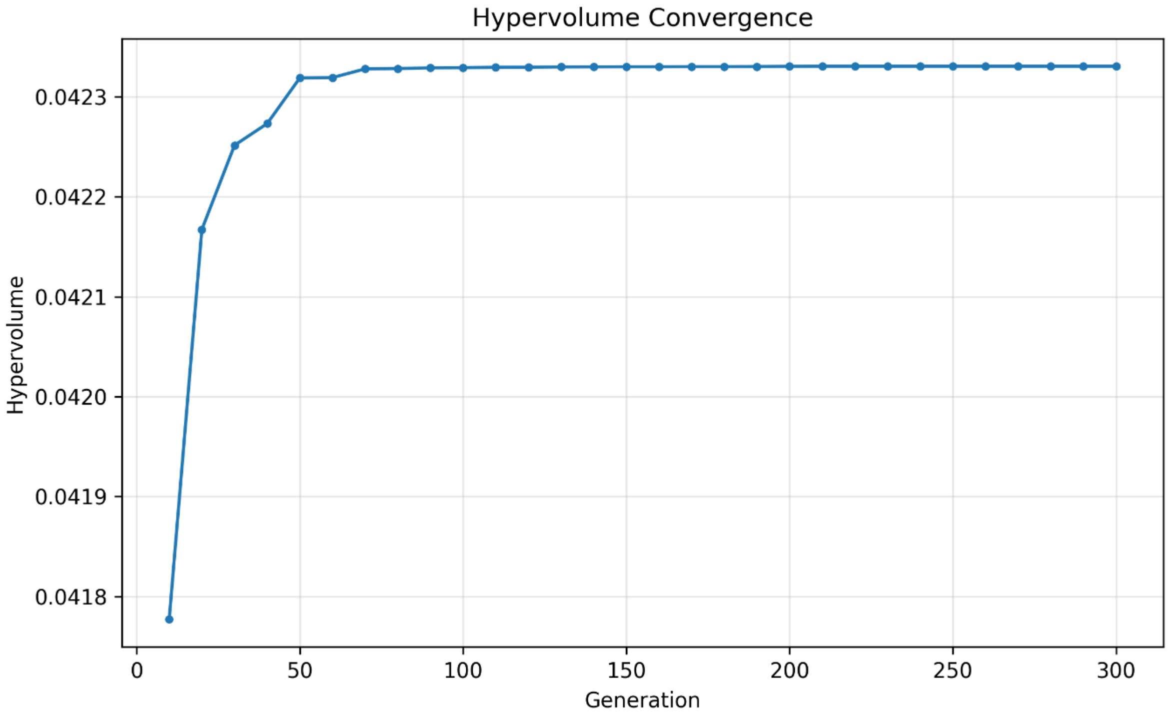

Figure 5-1 Hypervolume convergence of the reported LHS-initialised NSGA-II run.

The hypervolume increased rapidly during the early generations and approached a stable plateau by approximately generation 50 to 70, after which only marginal refinements were observed. At the reported numerical precision, the hypervolume remained unchanged from generation 210 onward, indicating stable objective-space coverage for the final nondominated archive.

Multi-seed stability was then assessed using five independent seeds: 11, 22, 33, 44, and 55. Table 5-2 reports the mean, standard deviation, minimum, maximum, and coefficient of variation for the main convergence and objective-space indicators. The multi-seed stability results further support the repeatability of the optimisation. The mean final hypervolume across the five independent seeds was 0.04233028, with a coefficient of variation of 0.00035%. This very small variation indicates that the objective-space boundary identified by the optimiser was stable across independent runs. The result does not prove mathematical global optimality, but it supports the use of the reported Pareto front as a reliable internal approximation within the fixed case-study model.

Table 5-2 Multi-seed stability .

Metric	Mean	Standard deviation	Minimum	Maximum	CV (%)
Final hypervolume	0.04233028	0.00000015	0.04233007	0.04233043	0.00035
IGD against multi-seed reference	0.00004612	0.00000313	0.00004290	0.00005065	6.78370
Final nondominated solutions	8,201.40	466.31	7,392.00	8,502.00	5.68576
Minimum normalised cost	0.98391274	0.00000019	0.98391260	0.98391307	0.00002
Minimum normalised carbon	0.97409071	0.00000060	0.97409033	0.97409176	0.00006
Minimum normalised time	0.93241225	0.00003403	0.93239064	0.93247238	0.00365
Normalised cost range	0.17187986	0.00003013	0.17182644	0.17189782	0.01753
Normalised carbon range	0.01513872	0.00003002	0.01510269	0.01517877	0.19830
Normalised time range	0.15011403	0.00006900	0.15000058	0.15017938	0.04597

The multi-seed results show that the final hypervolume was highly repeatable. The small variation in minimum normalised CFS framing process cost, CFS framing A1–A5 embodied carbon, and CFS framing process-hours also indicates that the main objective-space boundaries were consistently identified. The higher CV for IGD is acceptable in this context because IGD is sensitive to local point spacing and distribution along the nondominated front, while the hypervolume result confirms stable overall objective-space coverage.

The convergence and multi-seed results confirm optimisation repeatability within the implemented response model. They do not, by themselves, validate the physical accuracy of the cost, carbon, or process-hour equations. A stable hypervolume result means that the algorithm repeatedly recovered the same model-defined trade-off region; it does not prove that the model coefficients represent actual market prices, certified carbon values, or measured site productivity. For this reason, the Pareto front is interpreted together with the S2 comparison, component breakdown, sensitivity analysis, and Monte Carlo robustness results.

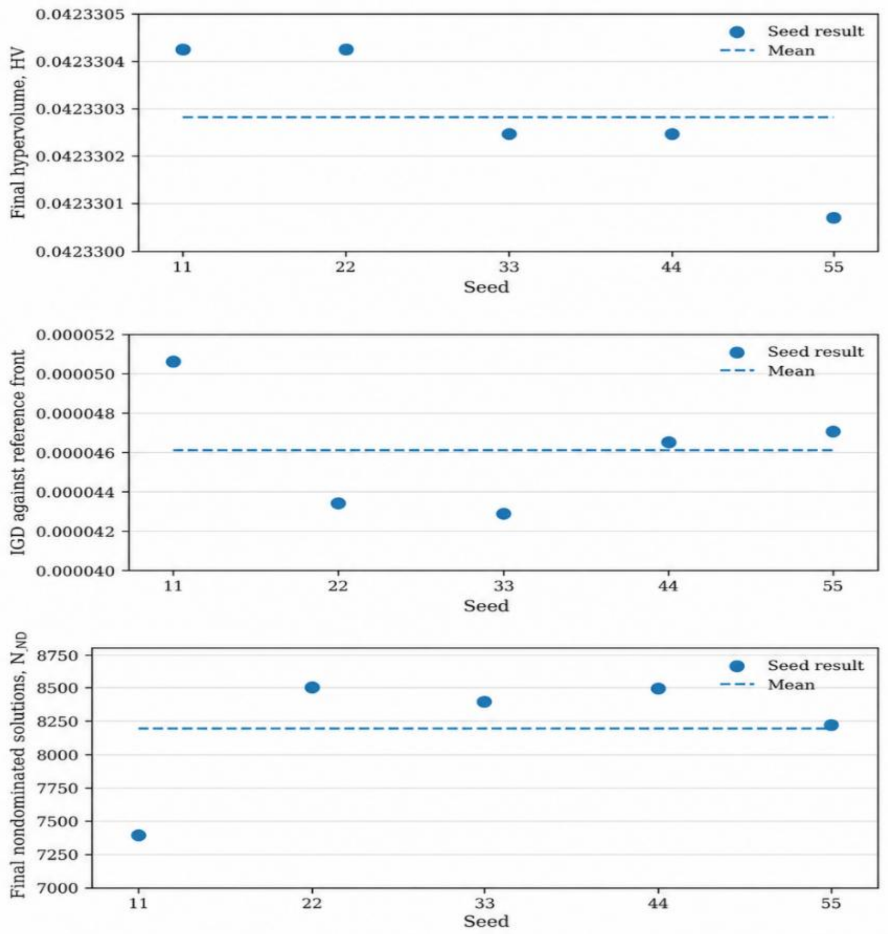

Figure 5-2 Multi-seed stability plot.

Figure 5-2 shows the seed-level stability of the final hypervolume, IGD, and nondominated archive size across the five independent NSGA-II runs. The final hypervolume values were closely grouped around the mean, confirming that the overall objective-space coverage was highly repeatable. The hypervolume axis is shown over a narrow numerical range to make the very small seed-to-seed differences visible. IGD showed greater variation than hypervolume, indicating that the spacing of nondominated points relative to the reference front was more sensitive to stochastic search variation than the outer objective-space coverage. The final nondominated-solution count also varied between seeds, but all runs produced a large nondominated archive. Overall, the seed-level results support the interpretation that the reported Pareto front is a stable NSGA-II approximation rather than a seed-specific outcome.

A grid-reference check was also used as an internal validation benchmark. The reported NSGA-II front contained 8,302 non-dominated solutions, while the grid-reference Pareto set contained 12 nondominated solutions. The NSGA-II approximation achieved an IGD against the grid-reference front of 0.00307. The NSGA-II front also identified slightly lower normalised extreme values than the grid reference for cost, carbon, and time. The minimum normalised CFS framing process cost was 0.98391 for NSGA-II compared with 0.98410 for the grid reference; the minimum normalised CFS framing A1–A5 embodied carbon was 0.97409 compared with 0.97418; and the minimum normalised CFS framing process-hours was 0.93241 compared with 0.93723. This does not prove global optimality, but it supports the conclusion that the NSGA-II search recovered a finer and more extensive approximation than the discrete grid-reference check used for internal validation.

Finally, a solution-count audit was used to verify that the reported Pareto-front size was not confusing with the total number of candidates evaluated. As shown in Table 5-3, the reported model run completed 47,549 evaluated fitness calls. All evaluated candidates satisfied the hard constructability filters, and no hard-infeasible evaluated candidates were recorded. The final external nondominated archive contained 8,302 solutions, and the same number remained after physical-objective filtering. Therefore, the Pareto-front approximation presented in the following sections are based on 8,302 feasible nondominated solutions.

Table 5-3 Solution count audit for the reported model run.

Metric	Value	Interpretation
Total evaluated fitness calls	47,549	Total candidate evaluations completed during the NSGA-II process
Hard-feasible evaluated candidates	47,549	Evaluated candidates satisfying the hard constructability constraints
Hard-infeasible evaluated candidates	0	No evaluated candidates violated the hard constraint filters
External nondominated archive size	8,302	Nondominated archive returned by the optimisation process
Final analysed nondominated Pareto solutions	8,302	Finite nondominated solutions retained after physical-objective filtering
Representative selected solutions	5	Selected solutions used for later comparison and decision-support reporting

The convergence and stability results support the use of the final nondominated archive as an internally stable approximation for the implemented response model; they do not imply recovery of an exact global Pareto set.

5.2: LHS Initial Population Coverage

Section 5.2 reports the coverage achieved by the implemented LHS initial population; the rationale for LHS was established in Chapter 2 and the implementation settings were defined in Chapter 3.

The generation-zero population was initialised using Latin Hypercube Sampling to improve coverage of the four optimiser-controlled variables before NSGA-II selection began. The reported optimisation used a population size of two hundred and evaluated the continuous design vector $\left(x = \left[R_p, R_t, OC, P_f\right]\right)$, where (R_p) is panel rationalisation intensity, (R_t) is truss rationalisation intensity, (OC) is opening/detailing complexity, and (P_f) is prefabrication level. The purpose of this stage was to reduce the risk of an initially clustered population and to ensure that each variable range was represented before crossover, mutation, nondominated sorting, and archive formation.

Table 5-4 LHS initial population audit for the reported optimisation run.

Variable	Minimum	25th percentile	Median	Mean	75th percentile	Maximum	Engineering interpretation
($\boldsymbol{R_p}$), panel rationalisation intensity	0.00369	0.24953	0.49983	0.49994	0.74814	0.99566	Near-full coverage of panel standardisation intensity
($\boldsymbol{R_t}$), truss rationalisation intensity	0.00210	0.25259	0.50039	0.50010	0.74943	0.99810	Near-full coverage of truss standardisation intensity
($\boldsymbol{OC}$), opening/ detailing complexity	0.80030	0.92626	1.04926	1.05004	1.17413	1.29788	Coverage from simplified to more complex opening/detailing cases
($\boldsymbol{P_f}$), prefabrication level	0.50186	0.60002	0.70015	0.69996	0.79875	0.89833	Coverage from moderate to high prefabrication levels

Table 5-4 shows that the initial two hundred design vectors provided near-complete coverage of the prescribed variable bounds. Both rationalisation variables extended from values close to zero to values close to one, indicating that the initial population included low, moderate, and high standardisation cases for both wall-panel and roof-truss rationalisation. The opening/detailing complexity variable covered the prescribed range from simplified detailing to higher-complexity detailing, while prefabrication level covered the range from lower to higher off-site production scenarios. This confirms that the optimisation did not begin from a narrow subset of the design space.

5.3: Final Pareto Front

The final Pareto front represents a structured decision set rather than a single optimum. Each nondominated solution reflects a different balance between CFS framing process cost, CFS framing A1–A5 embodied carbon, and CFS framing process-hours. Because all three objectives are minimised, movement along the front represents a change in engineering priority rather than a simple improvement in every output. The front therefore provides the basis for identifying cost-priority, carbon-priority, time-priority, balanced, and robustness-preferred alternatives.

The model-calculated S2 reference scenario was used as the benchmark for interpreting the Pareto front. S2 was defined by the reference vector [0.0, 0.0, 1.0, 0.72], with a CFS framing process cost of NZD 18,335.47, CFS framing A1–A5 embodied carbon of 5,366.79 kgCO_{2e}, and CFS framing process-hours of 93.92 h. These values are used only as the model-calculated comparison point for the fixed CFS framing case study; they are not measured commercial project data.

For practical interpretation, the term “Pareto front” is used later to refer to the interpreted trade-off surface, while “nondominated archive” refers to the full set of retained NSGA-II solutions.

Table 5-5 Objective ranges of the final Pareto front relative to the S2 reference scenario.

Objective	Unit	S2 reference	Pareto minimum	Pareto maximum	Best change vs S2	Pareto range vs S2: best	Engineering interpretation
CFS framing process cost	NZD	18,335.47	18,040.50	21,192.33	−1.61%	+15.58%	The front includes lower-cost options, but the lowest-carbon region is associated with higher process cost
CFS framing A1–A5 embodied carbon	$kgCO_{2e}$	5,366.79	5,227.74	5,309.08	−2.59%	−1.08%	All final nondominated solutions reduce embodied carbon relative to S2
CFS framing process-hours	Hr	93.92	87.57	101.67	−6.76%	+8.25%	The front includes lower-time options, but some low-carbon alternatives require longer process-hours

Note: S2-relative changes are signed percentages. Negative values indicate improvement relative to S2, while positive values indicate a penalty. The “least favourable” value is the weakest improvement or largest penalty within the final Pareto-front range, depending on the objective.

The improvement magnitudes in Table 5-5 are modest. The lowest-cost solution reduces CFS framing process cost by only 1.61% relative to S2, while the maximum carbon reduction across the front is 2.59% and the maximum process-hour reduction is 6.76%. These values are meaningful within the deterministic response model because they identify the direction and structure of the available improvement space. However, they are not large enough to be interpreted as guaranteed project savings under real construction conditions. In practice, supplier quotation differences, labour productivity variation, steel-price movement, detailing changes, transport assumptions, and carbon-factor uncertainty could be larger than the deterministic differences between some Pareto alternatives. The value of the Pareto front is therefore not a claim of large universal savings, but a transparent decision-support map showing which configurations are preferable within the stated model boundary and which alternatives require trade-offs.

The deterministic carbon spread across the final Pareto front is narrow. The nondominated carbon range is $5309.08 - 5227.74 = 81.34$ $kgCO_2e$, which is approximately $81.34/5366.79 \times 100 = 1.52\%$ of the S2 carbon reference. This means that deterministic carbon differences between neighbouring Pareto alternatives should not be over-interpreted as practically separable unless they remain distinguishable under the Monte Carlo uncertainty ranges reported in Chapter 6.

The cost and process-hour results show a more mixed pattern. The lowest-cost Pareto solution has a process cost of NZD 18,040.50, which is 1.61% lower than S2. The highest-cost Pareto solution has a process cost of NZD 21,192.33, which is 15.58% higher than S2. Similarly, the process-hours range from 87.57 h to 101.67 h, corresponding to a best reduction of 6.76% and a worst increase of 8.25% compared to S2. This confirms that the final Pareto front contains both direct improvement solutions and trade-off solutions.

Figure 5-3 shows that the nondominated archive contains two distinct decision regions. The first region contains 2,695 solutions that improve process cost, A1–A5 carbon, and process-hours simultaneously relative to S2. The second region contains 5,607 nondominated solutions that remain useful trade-off alternatives but do not improve all three objectives at the same time. This distinction is important because a nondominated solution is not automatically a practical improvement over the reference case.

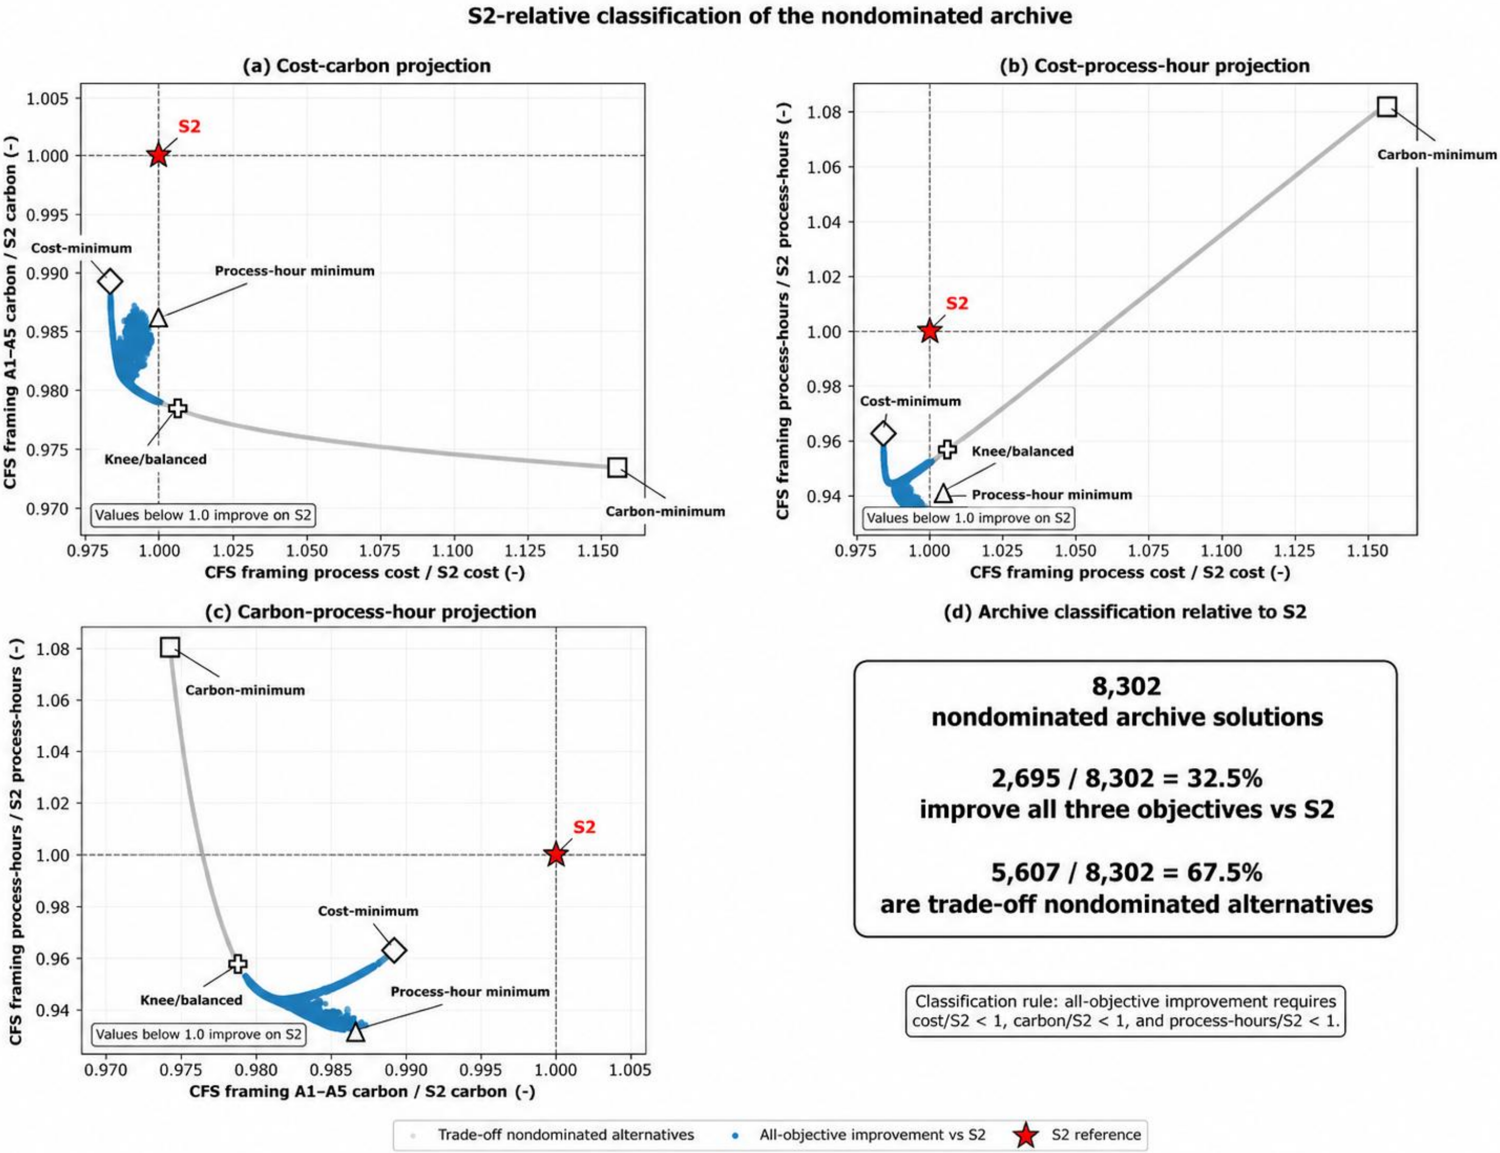

Figure 5-3 S2-relative classification of the nondominated archive.

The figure separates the 8,302 nondominated solutions into 2,695 configurations that improve CFS framing process cost, CFS framing A1–A5 carbon, and CFS framing process-hours simultaneously relative to S2, and 5,607 nondominated configurations that remain trade-off alternatives. Reference lines at 1.0 mark the S2 objective values; values below 1.0 indicate improvement relative to S2. The labelled representative points show the cost-minimum, carbon-minimum, process-hour-minimum, and knee/balanced solutions.

The presence of 2,695 all-objective improvement solutions should not be interpreted as evidence that the objectives are non-conflicting or that S2 was invalid. It indicates that the model-calculated S2 reference is not locally efficient under the implemented response model. The trade-off structure remains visible because the archive still contains carbon-priority, time-priority, cost-priority, and balanced alternatives with different cost, carbon, and process-hour consequences. The practical value of the optimisation is therefore the classification of improvement and trade-off regions, not the claim of large universal savings.

Table 5-6 Interpretation of nondominated archive classes relative to S2.

Archive class	Count	Meaning	Interpretation in this study
All-objective improvement solutions	2,695	Cost, A1–A5 carbon, and process-hours are all below S2	Candidate early-stage configurations that improve all three objectives under the implemented response model
Trade-off nondominated solutions	5,607	At least one objective improves, but not all three improve together	Useful priority-specific alternatives, especially for carbon-priority or prefabrication-priority decisions
Total nondominated archive	8,302	Solutions not dominated by another feasible solution	Stable decision set for S2-relative interpretation, not a claim of universal optimality

The classification in Table 5-6 is used throughout the chapter to prevent conflating Pareto membership with S2-relative improvement. This distinction is central to the decision-support interpretation of the archive.

A direct S2 comparison provides a clearer view of this distinction. Of the 8,302 final nondominated solutions, 2,695 solutions strictly improved all three objectives relative to S2. If equality with S2 is included, 2,696 solutions were equal to or better than S2 across all three objectives. All 8,302 solutions reduced CFS framing A1–A5 embodied carbon. For process cost, 2,696 solutions were equal to or below S2, while 5,606 solutions were above S2. For process-hours, 5,319 solutions were equal to or below S2, while 2,983 solutions were above S2. These counts show that the Pareto front includes a meaningful group of simultaneous improvement solutions, but also a larger set of carbon-focused trade-off solutions.

The size of the final nondominated archive should be interpreted in relation to the continuous response-model formulation. The optimiser does not select from a small discrete library of panel and truss schemes; it searches a four-variable continuous design space defined by R_p, R_t, OC, and P_f. In such a response surface, many solutions can be nondominated because movement along the front improves one objective while worsening, or only slightly changing, another. The 8,302 archive points therefore represent a dense approximation of the trade-off surface rather than 8,302 separate construction recommendations. Some points may be numerically distinct but practically close, which is why the thesis interprets the front through objective ranges, representative solutions, S2-relative classification, component breakdowns, TOPSIS ranking, and dashboard filtering rather than by discussing every point individually. The presence of 2,695 solutions that improve all three objectives relative to S2 should not be read as proof that S2 was a poor real-world design. S2 is a reference configuration, not an assumed Pareto-optimal design. The result indicates that, within the implemented response model, coordinated rationalisation, reduced detailing burden, and adjusted prefabrication level can produce modelled improvements relative to that reference condition.

Table 5-7 S2 comparison count for final nondominated solutions.

Comparison category	Number of solutions	Interpretation
Final nondominated Pareto solutions	8,302	Feasible solutions retained for final analysis
Strictly better than S2 in all three objectives	2,695	Direct simultaneous improvement in cost, carbon, and process-hours
Equal to or better than S2 in all three objectives	2,696	Includes solutions equal to S2 in one objective
Carbon equal to or below S2	8,302	All final nondominated solutions reduce CFS framing A1–A5 embodied carbon
Cost equal to or below S2	2,696	Lower-cost or cost-neutral subset
Process-hours equal to or below S2	5,319	Lower-time or time-neutral subset

The cost–process-hours projection does not show a single uniform relationship. Instead, the front includes a lower-cost/lower-time region and a higher-cost/higher-time branch. This is important from a construction perspective because it shows that time reduction is not simply obtained by spending more. Some configurations reduce process-hours while remaining close to or below the S2 cost benchmark, whereas other configurations accept higher cost and longer process-hours to obtain additional carbon reduction.

The carbon–process-hours projection shows the clearest trade-off. The lowest process-hour region occurs at a moderate carbon reduction, while the lowest-carbon region is associated with longer process-hours. Therefore, within this case study, the deepest carbon reduction is linked with increased process-hours and higher process cost. This is interpreted as a case-study-specific trade-off caused by the response model and fixed CFS framing boundary, not as a general rule for all CFS buildings.

The branch distribution shown in Figure 5-4 also indicates a stable design tendency within the reported case-study boundary. Most final nondominated solutions occur in the derived panel Scheme 2 / truss Scheme 2 branch, while only a small number of derived panel Scheme 3 / truss Scheme 2 solutions appear near the low-time region. This does not indicate a failure of the optimisation. Rather, it suggests that the response model consistently favoured a preferred standardisation branch for this specific CFS framing geometry, objective formulation, and constraint boundary. The detailed design-variable spread is discussed separately in the following section.

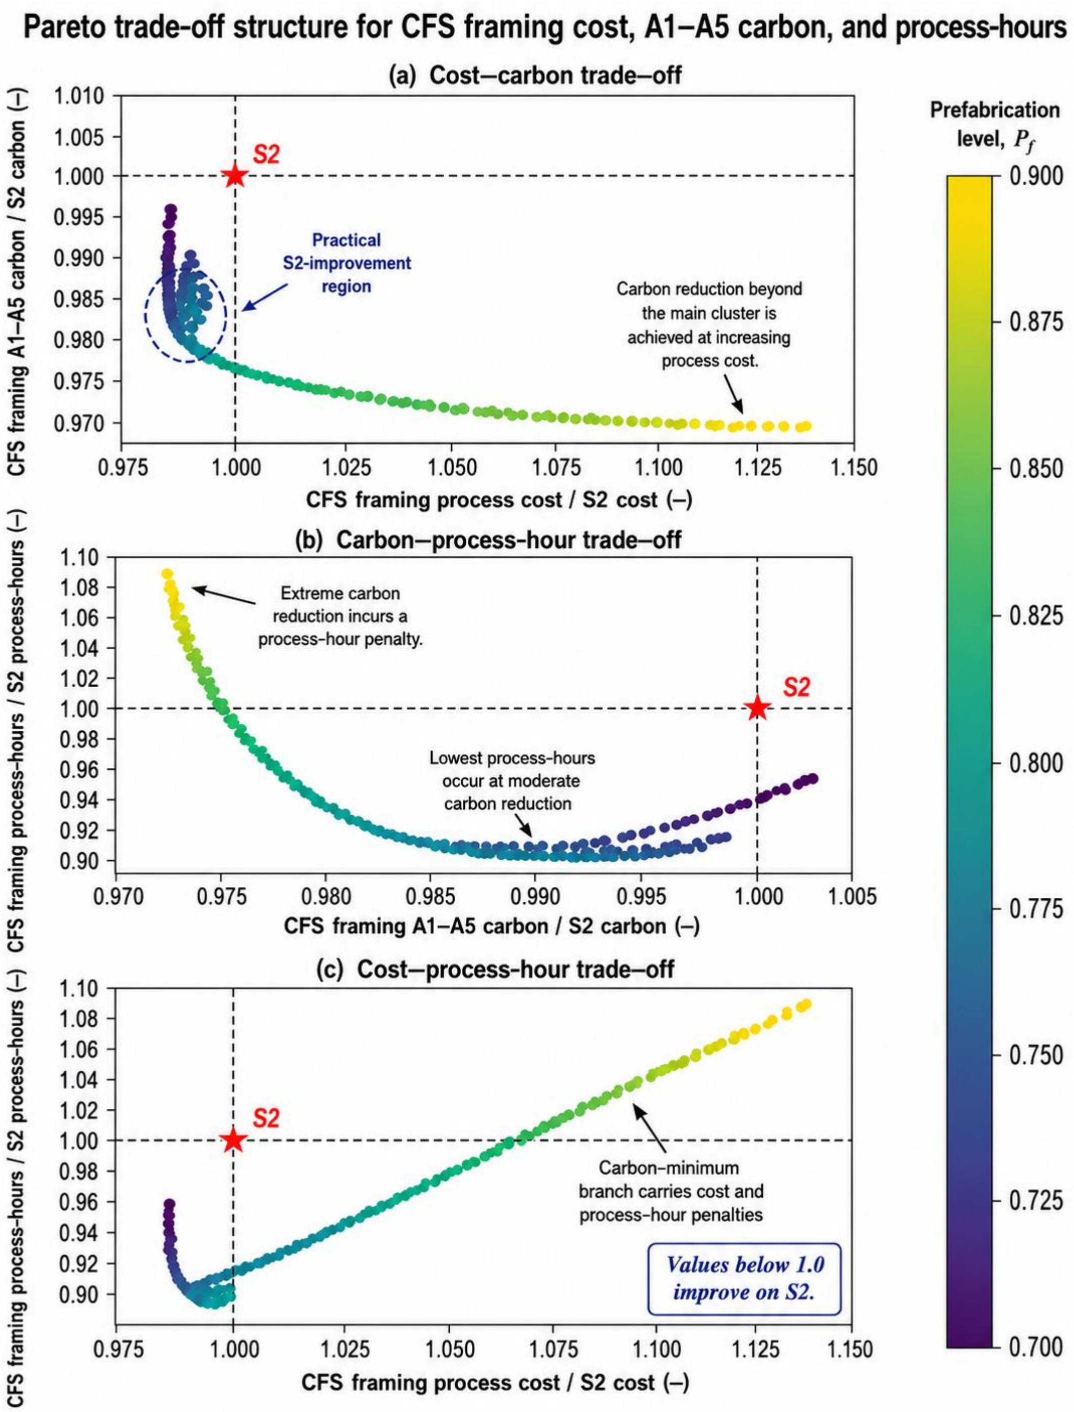

Figure 5-4 Two-dimensional projections of the nondominated archive with S2 reference and representative solutions.

Figure 5-4 shows that the main conflict in the implemented response model is not a simple cost-versus-carbon conflict. Many alternatives reduce carbon and cost together relative to S2. The sharper trade-off appears when carbon-priority configurations move toward higher prefabrication levels, where additional factory, logistics, handling, and coordination effects increase process cost and process-hours. The dominant engineering conflict is therefore between carbon reduction through prefabrication-intensive configurations and the process penalties associated with that prefabrication, rather than between carbon reduction and cost in isolation.

The final archive was interpreted using unique decision-vector and objective-value outputs generated by the reported optimisation run. Because the design variables are continuous, very small numerical differences may not correspond to materially different CFS framing decisions. For this reason, this study does not interpret all 8,302 archive points individually. Instead, it uses objective ranges, representative archetypes, S2-relative classification, component breakdowns, and dashboard filtering to translate the large archive into engineering-relevant alternatives.

5.4: NSGA-II Versus Grid Reference

Section 5.4 is retained because it provides a diagnostic comparison against a grid-reference approximation, addressing whether the NSGA-II archive is consistent with a direct design-space sampling check.

The grid-reference comparison provides an internal model-level check on the NSGA-II approximation. The grid front was generated from a finite grid over the bounded four-variable response space and is therefore not treated as the true global Pareto front. Instead, it provides an independent reference set against which the NSGA-II archive can be compared. Agreement between the NSGA-II front and the grid-reference front supports the credibility of the reported Pareto approximation, while the evolutionary search provides a denser and more diverse nondominated archive for interpretation.

The grid check evaluated 2,646 feasible CFS framing configuration cases. From these cases, twelve nondominated grid-reference solutions were obtained. By comparison, the reported NSGA-II run produced 8,302 feasible nondominated solutions. This difference is expected because the grid-reference set is a finite and comparatively coarse diagnostic sample, whereas NSGA-II searches the continuous bounded response space and retains a dense nondominated archive. The grid comparison is therefore used to check whether NSGA-II is recovering the same general trade-off region, not to match the archive size point-for-point.

Table 5-8 shows that the NSGA-II and grid-reference fronts are closely aligned in the main objective ranges. The NSGA-II front identified slightly lower minimum values for CFS framing process cost, CFS framing A1–A5 embodied carbon, and CFS framing process-hours. The differences are small in absolute engineering terms: approximately NZD 3.36 for minimum cost, 0.47 $kgCO_{2e}$ for minimum carbon, and 0.45 hr for minimum process-hours. This close agreement supports the credibility of the NSGA-II approximation, while the much larger NSGA-II solution count provides a more detailed decision-support set.

Table 5-8 NSGA-II versus grid-reference comparison.

Metric	NSGA-II Pareto front	Grid-reference front	Engineering interpretation
Total evaluated grid cases	—	2,646	Number of feasible grid cases used to form the reference comparison
Nondominated solutions	8,302	12	NSGA-II provides a denser set of decision alternatives

Metric	NSGA-II Pareto front	Grid-reference front	Engineering interpretation
Minimum normalised cost	0.98391	0.98410	NSGA-II identified a slightly lower-cost boundary
Maximum normalised cost	1.15581	1.15969	Both methods captured the higher-cost carbon-focused region
Minimum CFS framing process cost	NZD 18,040.50	NZD 18,043.86	Difference is small, supporting consistency between methods
Maximum CFS framing process cost	NZD 21,192.33	NZD 21,263.51	Grid reference extends slightly higher in cost
Minimum normalised carbon	0.97409	0.97418	NSGA-II identified a slightly lower-carbon boundary
Maximum normalised carbon	0.98925	0.99188	NSGA-II final front remains below S2 carbon across its reported range
Minimum CFS framing A1–A5 carbon	5,227.74 $kgCO_{2e}$	5,228.21 $kgCO_{2e}$	Difference is exceedingly small, indicating close agreement at the low-carbon boundary
Maximum CFS framing A1–A5 carbon	5,309.08 $kgCO_{2e}$	5,323.20 $kgCO_{2e}$	Grid reference includes a slightly higher carbon nondominated point
Minimum normalised process-hours	0.93241	0.93723	NSGA-II identified a lower process-hour boundary
Maximum normalised process-hours	1.08250	1.08320	Both fronts captured similar upper time range
Minimum CFS framing process-hours	87.57 hr	88.02 hr	NSGA-II identified a lower-time configuration
Maximum CFS framing process-hours	101.67 hr	101.73 hr	Upper process-hour values are closely aligned
IGD against grid-reference front	0.00307	0.00000	Small distance value supports NSGA-II approximation consistency

The inverted generational distance, reported as 0.00307, provides an additional numerical check against the grid-reference front. In this context, IGD is interpreted as a distance measure in the normalised objective space: a smaller value indicates that the NSGA-II front is close to the grid-derived reference front. The reported value is small, which supports the conclusion that the NSGA-II front is consistent with the grid-reference approximation. However, this remains an internal model-based check and should not be interpreted as proof of the exact global Pareto set.

Figure 5-5 compares the reported NSGA-II Pareto front with the grid-reference Pareto front in the normalised cost–carbon objective space. The model-calculated S2 benchmark is shown at the intersection of $\bar{C} = 1.0$ and $\bar{E} = 1.0$. In this figure, all plotted NSGA-II nondominated solutions lie below the S2 carbon reference line, confirming that the reported Pareto front improves CFS framing A1–A5 embodied carbon relative to S2 within the model boundary. The cost axis shows a wider spread: some solutions remain below the S2 cost benchmark, while others move to higher process cost in exchange for additional carbon reduction.

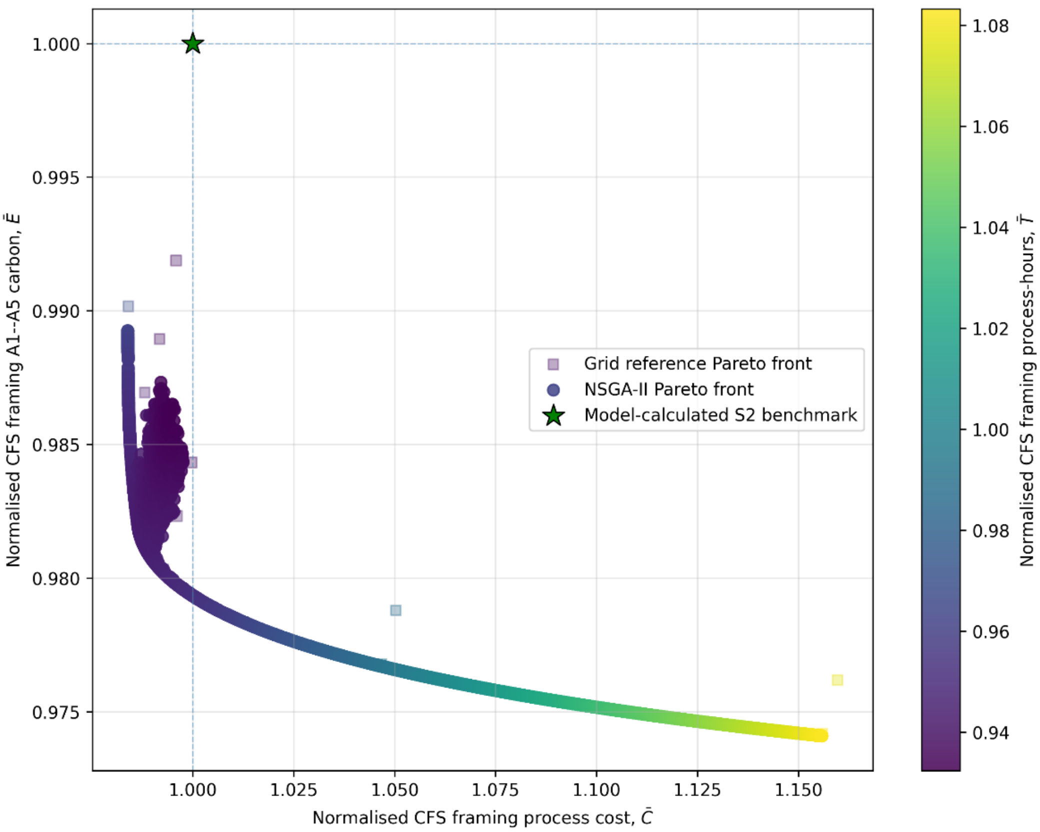

Figure 5-5 NSGA-II Pareto front compared with the grid-reference front in the normalised cost–carbon plane, with process-hours represented by colour.

The grid-reference front is represented by a small number of discrete points, whereas the NSGA-II front forms a much denser and more continuous set of nondominated alternatives. This supports the use of NSGA-II for decision support because it provides more intermediate CFS framing configuration options between the low-cost region and the lower-carbon region. The grid-reference points occupy the same general trade-off region, which provides internal validation that the NSGA-II approximation is consistent with the independent grid check.

The colour scale adds the third objective, normalised CFS framing process-hours. The darker points near the lower-cost cluster indicate lower process-hour solutions, while the lighter points toward the far right of the front indicate higher process-hours. This shows that, within the reported case study, the lowest-carbon region is associated with both higher cost and longer process-hours. Therefore, Figure 5-5 interpreted as evidence of a structured cost–carbon–time trade-off, not as evidence that every Pareto solution improves every objective simultaneously.

Overall, the figures support the numerical comparison reported in Table 5-8. The NSGA-II front is consistent with the grid-reference front but provides a finer set of feasible CFS framing alternatives for practical selection. This comparison strengthens confidence in the reported Pareto approximation while remaining an internal model-based validation rather than proof of the exact global Pareto set.

The grid-reference comparison supports the use of NSGA-II as a practical Pareto-search method for the implemented continuous response model, while still treating the final archive as an approximation rather than an exact global Pareto set.

5.5: Pareto Archetypes and Representative Solutions

Representative solutions are used to translate the large Pareto archive into a small set of interpretable engineering decision positions. They do not separate optimisation runs and do not replace the full Pareto front. Instead, they identify selected nondominated alternatives that represent cost-priority, carbon-priority, time-priority, and balanced decision positions within the final archive.

Table 5-9 Representative Pareto solutions relative to the S2 reference scenario.

Representative	Cost-min	Carbon-min	Time-min	Knee/balanced
$\boldsymbol{R_p}$	0.6895	0.6895	0.7529	0.6907
$\boldsymbol{R_t}$	0.6785	0.6811	0.7217	0.6799
Panel scheme	2	2	3	2
Truss scheme	2	2	2	2
$\boldsymbol{OC}$	1	1	0.8	1
$\boldsymbol{P_f}$	0.7083	0.9	0.7825	0.8184
Cost (NZD)	18040.5	21192.24	18243.14	18446.19
Carbon ($kgCO_{2e}$)	5309.08	5227.74	5294.35	5253.04
Time (h)	90.45	101.67	87.57	89.94
Cost vs S2	−1.61%	+15.58%	−0.50%	+0.60%
Carbon vs S2	−1.08%	−2.59%	−1.35%	−2.12%
Time vs S2	−3.70%	+8.25%	−6.76%	−4.23%
Interpretation	Lowest-cost representative; improves all three objectives	Lowest-carbon representative; accepts cost and time penalties	Lowest process-hour representative; improves all three objectives	Balanced compromise with small cost premium

The representative solutions are selected to translate the large nondominated archive into interpretable decision positions. They are not claimed to be universally optimal designs. Each representative illustrates a different stakeholder priority within the implemented cost–carbon–process-hour response model.

Table 5-9 shows that the representative solutions describe different decision positions within the Pareto front. The Cost-min solution provides the lowest CFS framing process cost and reduces embodied carbon and process-hours relative to S2. The Time-min solution gives the largest process-hour reduction while also remaining slightly below S2 in cost and carbon. These two representatives therefore show that simultaneous improvement in all three objectives is possible within the reported case-study boundary.

The Carbon-min solution gives the greatest deterministic reduction in CFS framing A1–A5 embodied carbon, but the improvement is modest in magnitude and is achieved with clear process-cost and process-hour penalties. It should therefore be interpreted as a deterministic carbon-priority trade-off within the response model, not as generally superior or uncertainty-robust within the tested uncertainty ranges. Its practical attractiveness depends on whether the decision-maker is willing to accept a higher cost and longer process-hours for a small, modelled carbon reduction.

The deterministic representative-solution differences are read together with the uncertainty analysis in Chapter 6. In particular, the carbon differences among representative solutions are small relative to the carbon uncertainty propagated later through the Monte Carlo analysis. Therefore, the carbon-minimum representative is useful as an objective-space extreme, but it should not be treated as a statistically certain carbon improvement over nearby alternatives under uncertain carbon factors. The decision-support value of the representative set is that it exposes the trade-off structure clearly: cost-minimum and time-minimum solutions improve all three objectives relative to S2 within the deterministic model, while the carbon-minimum solution shows the cost and process-hour penalty required to reach the lowest modelled carbon value

Figure 5-6 Parallel-coordinated comparison of representative Pareto solutions. The grey lines represent the final nondominated Pareto solutions, while the coloured lines show the selected representative solutions. The figure compares rationalisation settings, opening/detailing complexity, prefabrication level, and normalised objective outcomes.

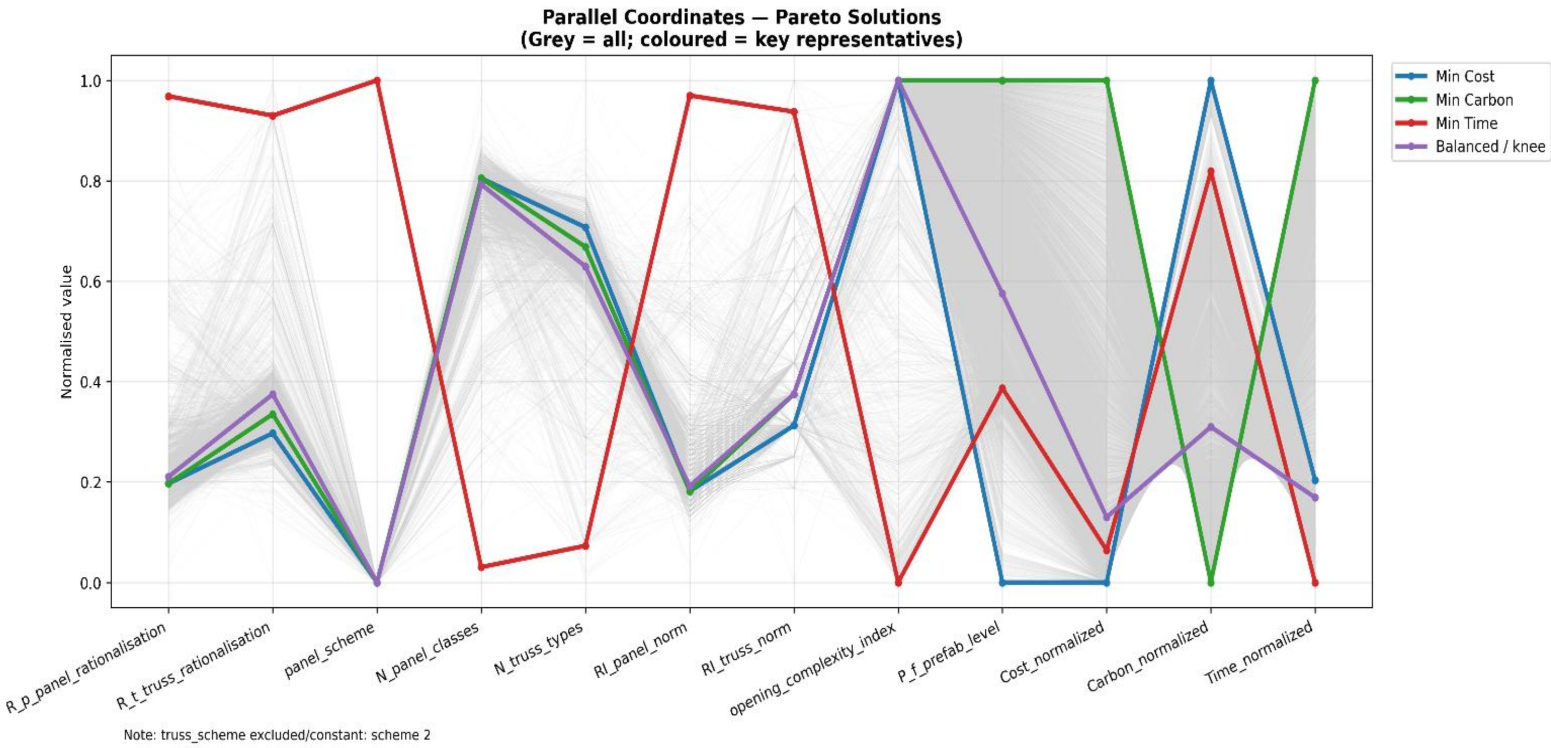

Figure 5-6 Parallel-coordinate comparison of representative Pareto solutions.

The grey lines represent the final nondominated Pareto-front solutions, while the highlighted lines show the selected representative alternatives. The figure compares rationalisation intensity, opening/detailing complexity, prefabrication level, *and normalised objective values, allowing the representative decision positions to be interpreted relative to the wider Pareto archive.*

Figure 5-6 shows that the representative solutions occupy different regions of the decision space. The cost-minimum and carbon-minimum representatives have similar rationalisation settings but differ mainly in prefabrication level and objective outcome. The time-minimum representative uses stronger panel rationalisation and lower opening/detailing complexity, which is consistent with its lower process-hours. The knee/balanced solution sits between the extreme cases and provides a practical compromise for projects where cost, embodied carbon, and process-hours must be considered together.

Overall, the representative-solution analysis converts the Pareto front into a small set of practical CFS framing decision options. The cost- and time-focused representatives show direct improvement relative to S2, the carbon-focused representative shows the trade-off required to reach the lowest carbon value, and the balanced representative provides a moderate alternative for multi-criteria decision-making.

5.6: Derived Panel and Truss Reporting-Category Distribution

Panel and truss schemes are interpreted here as derived reporting categories from the continuous rationalisation variables $R_p$and R_t. They are not independent optimiser variables. This distinction prevents the results from being read as a discrete scheme-selection optimisation when the implemented decision vector is $\mathbf{x} = [R_p, R_t, OC, P_f]$.

The derived scheme distribution should be interpreted as a reporting-level summary of the continuous rationalisation variables, not as evidence that the optimiser directly selected discrete fabrication schemes. The final Pareto front was concentrated mainly in Panel Scheme 2 and Truss Scheme 2, with a small Panel Scheme 3 branch appearing in the low-process-hour region. This indicates that the response model favoured a controlled standardisation band rather than maximum rationalisation across all variables.

Table 5-10 Panel and truss scheme distribution in the final Pareto front.

Panel scheme	Truss scheme	Pareto solutions	Share of Pareto front	Cost range (NZD)	Carbon range ($kgCO_{2e}$)	Time range (hr)	Interpretation
2	2	8,298	99.95%	18,040.50–21,192.33	5,227.74–5,309.08	87.58–101.67	Dominant case-study standardisation branch
3	2	4	0.05%	18,243.14–18,253.14	5,293.27–5,294.35	87.57–87.58	Small low-time branch with stronger panel rationalisation

The truss scheme's result is especially clear: all retained nondominated solutions occur in derived truss Scheme 2. This is interpreted as a case-specific response-model outcome rather than evidence that other truss schemes are generally inferior. Mechanistically, the retained

$R_t$band falls just above the rationalisation knee, where the model still receives repetition and setup benefits from reducing truss-family variation, but does not yet accumulate the stronger oversizing, grouping, handling, logistics, and setup penalties associated with more aggressive truss rationalisation. In other words, Scheme 2 survives nondominated sorting because it occupies the response-model balance point between insufficient truss repetition and excessive span-family grouping.

The panel scheme result is slightly more varied. While panel Scheme 2 dominates the final front, panel Scheme 3 appears in a very narrow region associated with the lowest process-hours. This aligns with the representative-solution analysis, where the time-minimum solution used panel Scheme 3, whereas the cost-minimum, carbon-minimum, and balanced solutions remained in panel Scheme 2. The result suggests that stronger panel rationalisation may be advantageous when schedule performance is prioritised, but the broader Pareto front is governed by panel Scheme 2.

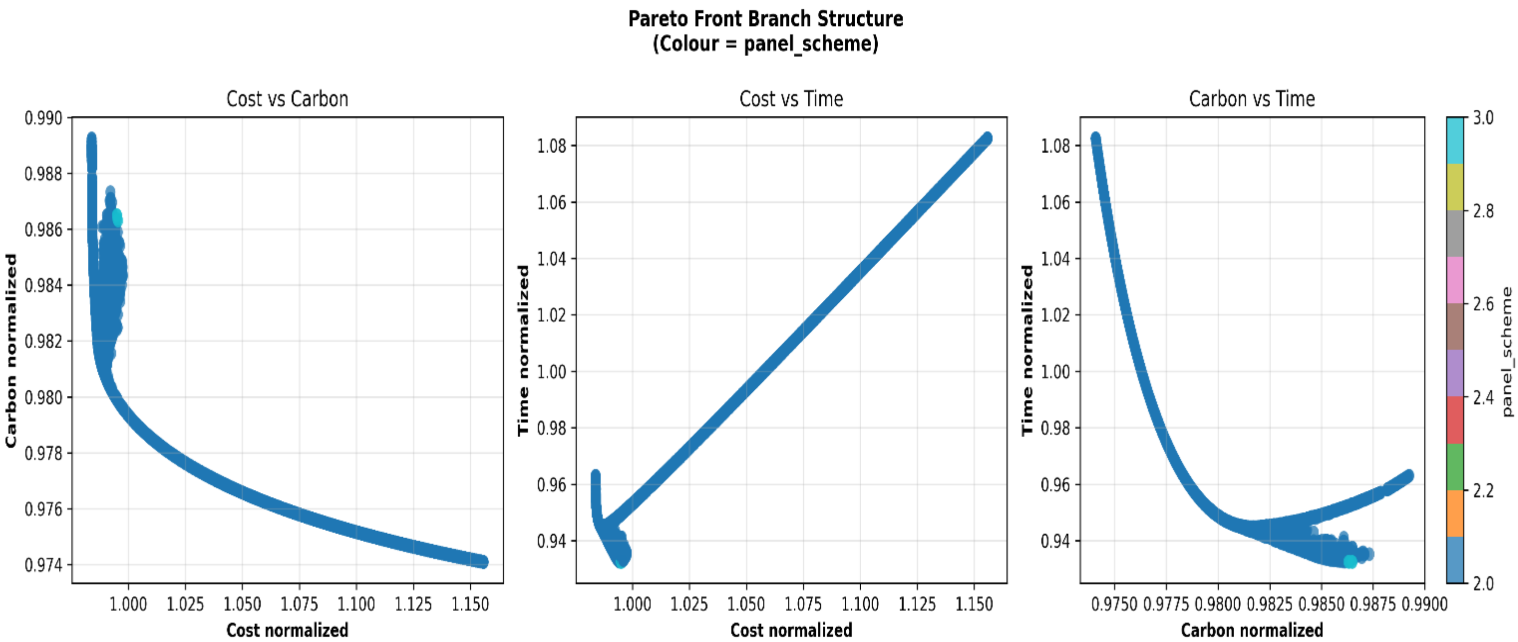

Figure 5-7 Distribution of derived panel-scheme categories across the final Pareto front in the normalised cost–carbon objective plan.

Figure 5-7 supports the distribution reported in Table 5-10. The plotted Pareto front is almost entirely represented by panel Scheme 2, with only a small set of panel Scheme 3 points visible near the low-time region. The cost–carbon projection shows that the dominant panel Scheme 2 branch covers almost the full trade-off surface, including both lower-cost and lower-carbon regions. The cost–time and carbon–time projections show that the small Scheme 3 group is associated with the shortest process-hour region rather than with the full cost–carbon range.

The ranges of R_p and R_t on the final Pareto front show that the nondominated solutions occupy a restricted rationalisation band rather than the full exploratory design space. This described as a stable design tendency within the case study, not as a collapse of design diversity. The optimisation clearly screened out less favourable scheme combinations and retained the branch that best balanced the competing objectives.

Overall, the discrete scheme distribution indicates that the final Pareto front is governed by a moderate-to-strong rationalisation strategy. For this case study, truss Scheme 2 is consistently retained, while panel Scheme 2 forms the principal decision branch and panel Scheme 3 appears only as a narrow low-time alternative. This narrows the practical decision space and suggests that design attention can be focused primarily on the Scheme 2 branch, with Scheme 3 considered only when process-hours are the dominant priority.

5.6.1: Derived Repetition-Index Interpretation

The retained Pareto front was further examined using the rationalisation-intensity variables R_p and R_t. In the optimisation formulation, R_p represents panel rationalisation intensity and R_t represents truss rationalisation intensity. The derived categories panel scheme and truss scheme are used only for reporting and interpretation. This section therefore examines how the retained R_p and R_t bands relate to the panel and truss scheme patterns observed in the final nondominated solution set.

The final Pareto front did not retain the full initial range of either rationalisation variable. Instead, the nondominated solutions concentrated within a moderate-to-strong rationalisation band. The retained R_p range was 0.6733 to 0.7555, with a mean of 0.6915, while the retained R_t range was 0.6582 to 0.7265, with a mean of 0.6817. This indicates that the final CFS framing alternatives favoured neither very low rationalisation nor maximum rationalisation. Rather, the best trade-off region was located within an intermediate-to-strong standardisation range.

Table 5-11 Rationalisation-intensity interpretation by retained Pareto branch.

Pareto branch	**R_p tendency**	**R_t tendency**	**Derived panel scheme**	**Derived truss scheme**	**Objective tendency**	**Engineering interpretation**
Dominant branch	0.6733–0.7555, mostly around 0.69	0.6582–0.7265, mostly around 0.68	Mainly Scheme 2	Scheme 2	Main cost–carbon–time compromise	Moderate-to-strong rationalisation provides the principal retained standardisation branch
Low-time branch	Near the upper (R_p) band	Within retained (R_t) band	Scheme 3	Scheme 2	Lowest process-hour region	Stronger panel rationalisation appears only in a narrow low-time region
Truss branch	Narrow (R_t) range	Scheme 2 only	Scheme 2 or 3	Scheme 2	Stable across retained solutions	Truss rationalisation shows a stable case-study preference rather than broad branch variation

The constant appearance of truss Scheme 2 is therefore not treated as an independently validated fabrication rule. It is a consequence of the implemented response equations, the S2 truss-family structure, the fixed steel-mass boundary, and the selected penalty coefficients for truss oversizing, setup, logistics, and site handling. The result is useful because it narrows the case-study decision space, but it should be recalibrated before being transferred to projects with different roof geometry, truss span distribution, transport constraints, or fabricator-specific truss libraries

The derived scheme results support this interpretation. Most Pareto solutions occur in panel Scheme 2 and truss Scheme 2, while only a small number of panel Scheme 3 solutions occur in the low-process-hour region. The truss category is constant at Scheme 2 across the final Pareto front. This should not be described as a modelling failure. It indicates that, for this case-study geometry and response model, truss Scheme 2 provided the most favourable retained balance between repetition benefit, span-family rationalisation, fabrication effort, embodied carbon, and process-hours.

Figure 5-8 Objective-response heatmaps for panel rationalisation intensity $R_p$and truss rationalisation intensity R_t. The horizontal and vertical axes represent the rationalisation-intensity variables, while the cell colour and annotated values show the best normalised objective response obtained after varying opening/detailing complexity and prefabrication level. Lower values indicate better performance because all three objectives are minimised. The heatmaps show that rationalisation benefits are not linear across the full range and are moderated by oversizing, trimming, handling, setup, and logistics penalties.

A clear pattern is visible in the heatmaps. Moving from extremely low rationalisation toward the moderate-to-strong range improves all three objectives. In other words, better repetition and standardisation tend to reduce process cost, reduce embodied carbon, and shorten process-hours up to a certain level. However, this improvement does not continue indefinitely. At the extreme upper end of the rationalisation scale, especially near $R_p = 1.0$ and $R_t = 1.0$, the objective values worsen. For example, the heatmaps show values of approximately 1.020 for cost, 0.999 for carbon, and 0.958 for process-hours at the most extreme corner. This indicates that maximum rationalisation is not the best overall strategy in this case study.

This result is important in engineering terms. It suggests that rationalisation is beneficial only within an appropriate range. Moderate-to-strong rationalisation improves repetition, fabrication efficiency, and installation productivity, but excessive rationalisation does not continue to improve performance. At high rationalisation intensity, the imposed standardisation can begin to create penalties through less efficient geometric matching, reduced flexibility, span grouping effects, or other constructability constraints embedded in the response model. The optimisation therefore supports a balanced interpretation: standardisation is valuable, but more standardisation is not automatically better.

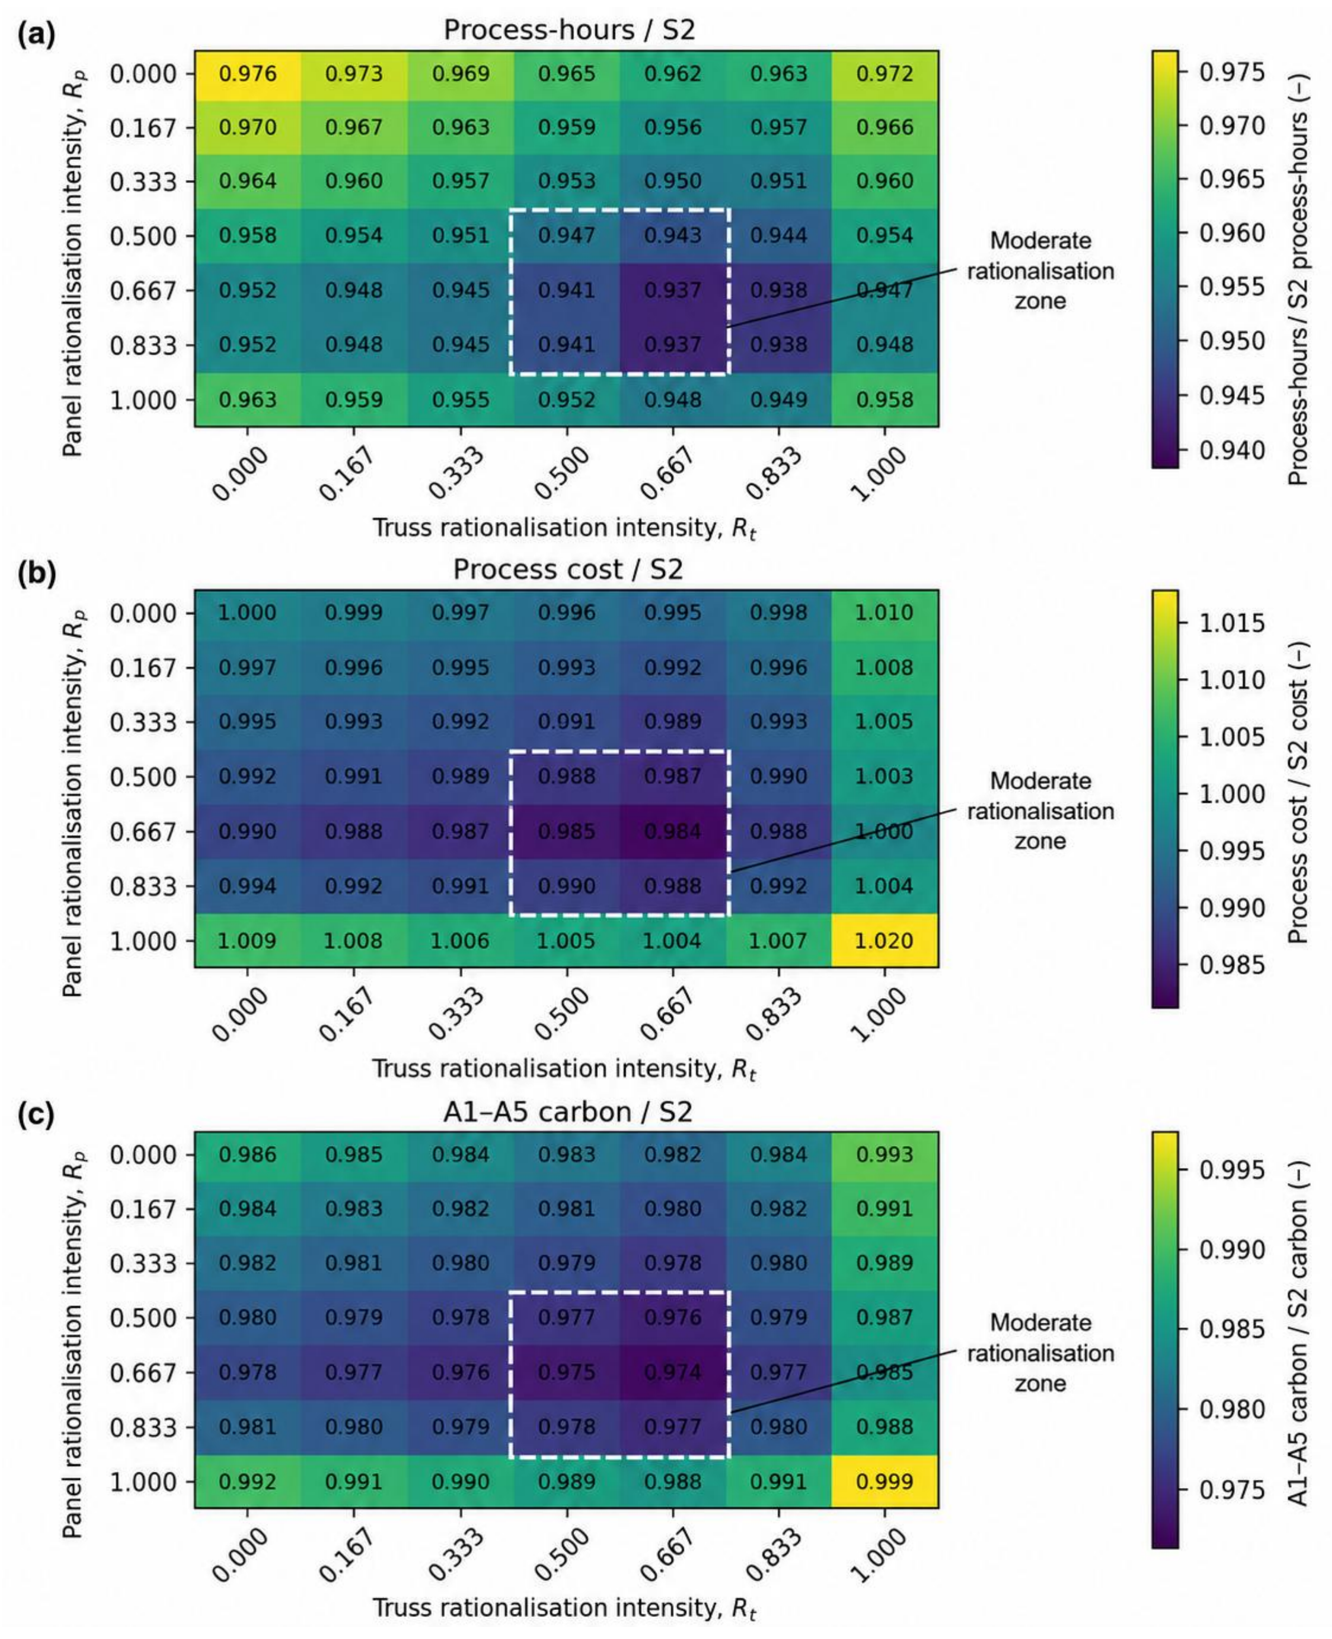

Figure 5-8 Rationalisation response heatmaps showing the moderate rationalisation zone

The heatmaps show S2-normalised process-hours, process cost, and A1–A5 carbon across panel rationalisation intensity $R_p$and truss rationalisation intensity R_t. The dashed box identifies the approximate moderate-rationalisation region $(R_p \approx 0.5\text{–}0.8,\ R_t \approx 0.5\text{–}0.8)$, where the implemented response model produces the most favourable combined objective behaviour. Values below 1.0 indicate improvement relative to S2. The annotated zone in Figure 5-8 shows that the most favourable response is not located at the maximum rationalisation boundary. Instead, the implemented model favours moderate panel and truss rationalisation, where repetition benefits are retained without incurring the full oversizing, grouping, or coordination penalties represented in the response model. This supports the interpretation that rationalisation should be controlled rather than simply maximised.

The relative influence of $R_p$and $R_t$is also interpretable. The retained Pareto front shows a stable case-study preference for truss Scheme 2, which indicates that truss rationalisation is important but does not produce multiple surviving truss branches in the final trade-off set. Panel rationalisation shows slightly more visible branch behaviour, because the small low-time branch occurs at stronger panel rationalisation levels and appears as panel Scheme 3. This means that panel rationalisation contributes more clearly to the distinction between the main compromise region and the narrow low-time region.

Overall, the retained $R_p$and $R_t$bands help explain the structure of the final Pareto front. The dominant nondominated solutions are not associated with either minimal or maximal rationalisation, but with a moderate-to-strong rationalisation band that provides the most favourable balance of process cost, A1–A5 embodied carbon, and process-hours. This finding is useful for CFS framing decision support because it narrows the practical design discussion to a preferred standardisation range rather than suggesting that the rationalisation variables should simply be driven to their extremes.

5.7: S2 Comparison and Component Breakdown Against S2

The final Pareto front was compared directly with the model-calculated S2 reference scenario to identify which solutions provide modelled improvement and which represent trade-offs. This check is necessary because a nondominated solution is not automatically better than S2 in every objective. A solution may reduce CFS framing A1–A5 embodied carbon while increasing process cost or process-hours. The S2 comparison therefore provides a transparent benchmark check before the representative solutions are used for decision-support interpretation.

Table 5-12 All-objective improvement and trade-off classification relative to S2.

Comparison category	Number of Pareto solutions	Share of Pareto front	Interpretation
Final nondominated Pareto solutions	8,302	100.0%	Feasible nondominated solutions retained for final analysis
Better than S2 in all three objectives	2,695	32.5%	Solutions that reduce cost, carbon, and process-hours simultaneously
Equal to or better than S2 in all three objectives	2,695	32.5%	Same as strict improvement in the reported output
CFS framing process cost equal to or below S2	2,695	32.5%	Cost-improving subset

Comparison category	Number of Pareto solutions	Share of Pareto front	Interpretation
CFS framing A1–A5 embodied carbon equal to or below S2	8,302	100.0%	All final Pareto solutions reduce embodied carbon relative to S2
CFS framing process-hours equal to or below S2	5,319	64.1%	Majority of solutions improve or match process-hours
CFS framing process cost above S2	5,607	67.5%	Carbon-focused trade-off region includes higher-cost solutions
CFS framing process-hours above S2	2,983	35.9%	Some lower-carbon solutions require longer process-hours

Table 5-12 shows that the final Pareto front contains a strict S2-improvement subset. A total of 2,695 solutions reduce all three minimisation objectives relative to S2 and can therefore be described as S2-dominating solutions under the definition established in Chapter 3. However, the whole Pareto front should not be described as uniformly superior to S2. All 8,302 solutions reduce embodied carbon relative to S2, but 5,607 solutions increase process cost, and 2,983 solutions increase process-hours. The final archive therefore contains both direct improvement solutions and deliberate trade-off solutions.

The modest magnitude of the S2-relative improvements does not make the optimisation result unimportant. Instead, it changes the interpretation of the contribution. For this fixed CFS framing case study, the feasible improvement space is narrow, meaning that large cost, carbon, and process-hour reductions should not be claimed. The decision-support value lies in identifying which configurations remain preferable within that narrow space, distinguishing direct S2-improvement solutions from trade-off solutions, and preventing decision-makers from assuming that the lowest-carbon solution is automatically the best practical option. In this sense, the optimisation framework is valuable because it makes a small and uncertain decision space explicit, rather than because it produces large headline savings.

The component breakdown is presented in delta-from-S2 form because the decision question is not only which representative has the lowest total value, but which process components change relative to the reference case. This format makes the trade-off mechanisms more visible than stacked total bars.

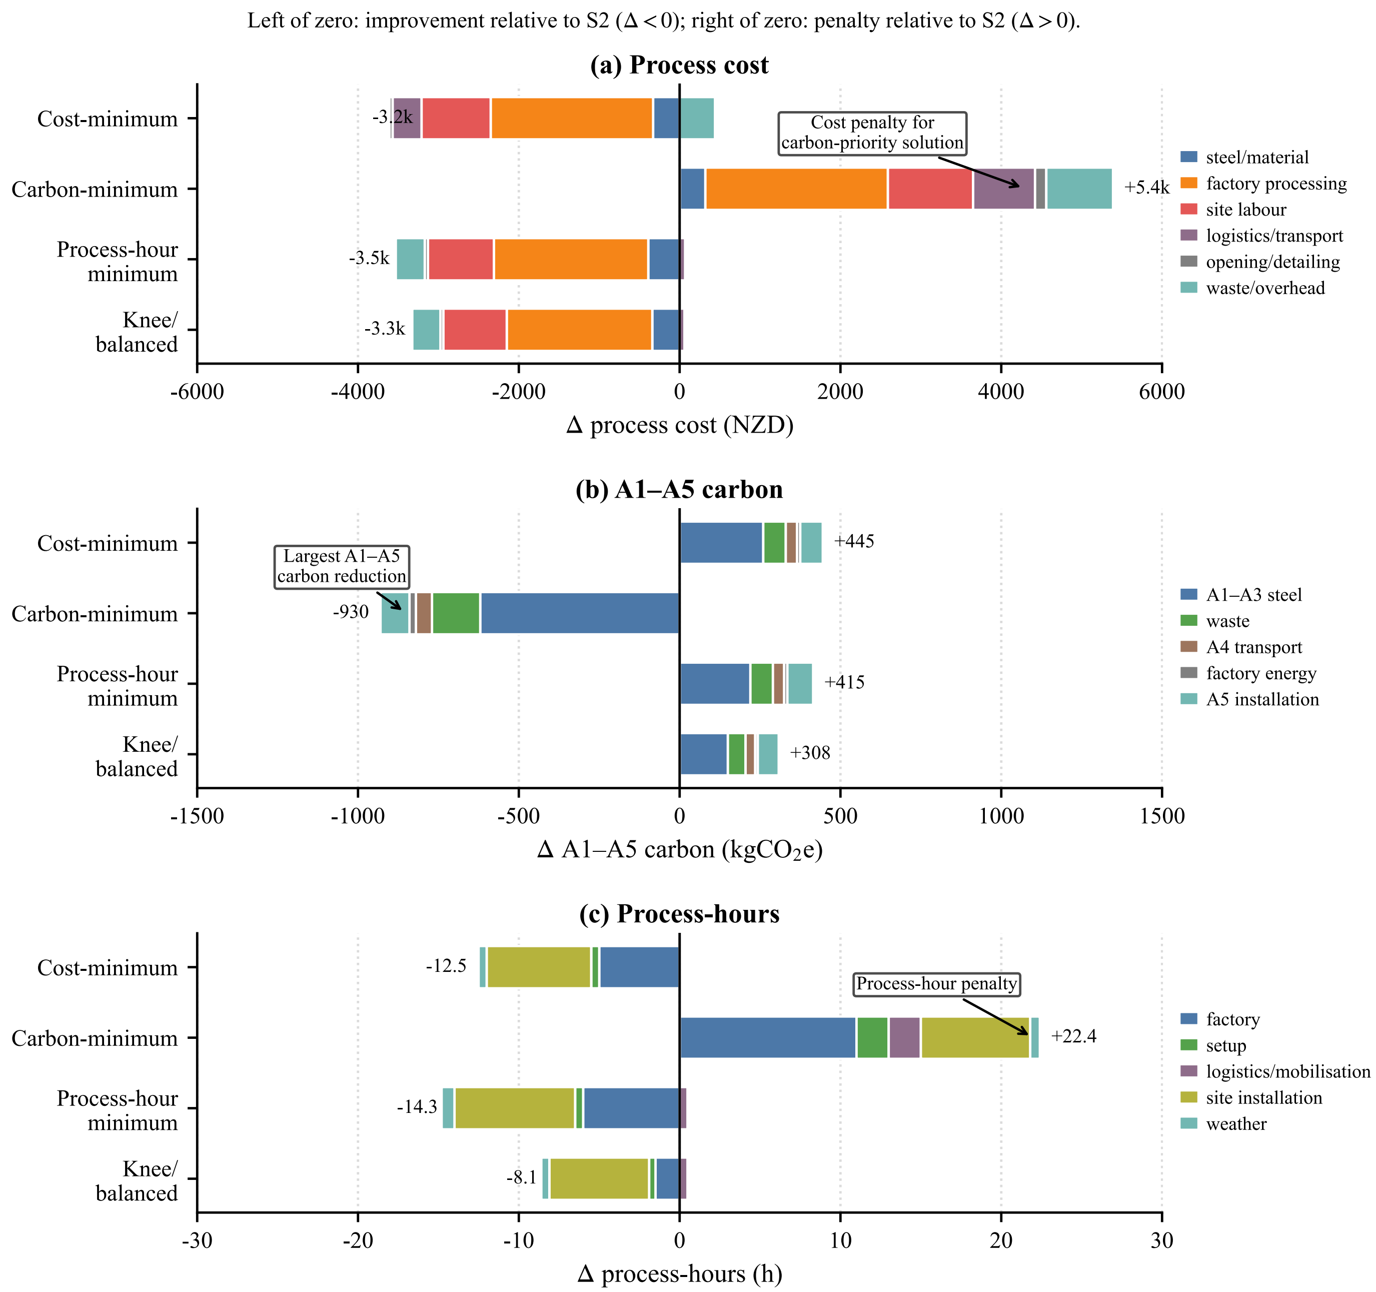

Figure 5-9 Component-level changes of representative solutions relative to S2

The figure shows process-cost, A1–A5 carbon, and process-hour component changes for the representative solutions relative to the S2 baseline. The zero line represents S2, negative values indicate improvement, and positive values indicate penalty. Presenting component changes rather than total stacked values makes the source of each trade-off visible, including the cost and process-hour penalty associated with the carbon-minimum representative.

Figure 5-9 shows that the representative solutions achieve their objective positions through different component-level mechanisms. The cost-minimum and process-hour-minimum representatives provide modest all-objective improvements, while the carbon-minimum representative reduces deterministic carbon at the expense of higher process cost and process-hours. This supports the interpretation that the carbon-minimum solution is a priority-specific trade-off rather than a generally preferred practical solution.

The component breakdown confirms that the objective improvements are not caused by one isolated model term. Cost reductions are mainly associated with steel mass response, factory processing, setup, site labour, logistics, waste, and overhead interactions. Carbon reductions are governed primarily by steel mass and A1–A3 steel carbon, with smaller contributions from A4 transport, factory energy, A5 installation, and waste allowance. Process-hour reductions depend on the redistribution of effort between factory production, setup, logistics, mobilisation, and site installation. This supports the interpretation that the optimisation result is a CFS production-configuration trade-off rather than a simple maximum-prefabrication or minimum-mass result.

Table 5-13 Component breakdown values for representative Pareto solutions relative to S2.

Representative	Total cost (NZD)	Cost vs S2	Total carbon ($kgCO_{2e}$)	Carbon vs S2	Total time (h)	Time vs S2	Main mechanism
Cost-min	18,040.50	−1.61%	5,309.09	−1.08%	90.45	−3.69%	Lower setup and site labour offset a slight increase in factory processing
Carbon-min	21,192.30	+15.58%	5,227.74	−2.59%	101.67	+8.25%	Lower steel and A5 carbon, but much higher factory processing cost and factory time
Time-min	18,243.19	−0.50%	5,294.40	−1.35%	87.57	−6.76%	Lower site installation time and lower site labour dominate the response
Knee/balanced	18,445.94	+0.60%	5,253.04	−2.12%	89.94	−4.24%	Moderate cost premium gives stronger carbon and time reduction than Cost-min
S2 reference	18,335.47	0.00%	5,366.79	0.00%	93.92	0.00%	Model-calculated benchmark

The cost breakdown shows that the Cost-min solution reduces total costs by approximately NZD 294.97 relative to S2. This occurs mainly because setup costs and site labour are lower, even though factory processing costs are slightly higher. The Carbon-min solution behaves differently. It has the lowest embodied carbon, but its total cost increases by approximately NZD 2,856.83 relative to S2, mainly because factory processing and overhead increase substantially. This explains why the carbon-minimum point is a trade-off solution rather than a general-purpose improvement.

The carbon breakdown shows that reductions are mainly associated with lower A1–A3 steel carbon and lower A5 installation carbon. The Carbon-min solution reduces total carbon by approximately 139.05 $kgCO_{2e}$ relative to S2. The Knee/balanced solution also performs strongly, reducing carbon by approximately 113.75 $kgCO_{2e}$, while avoiding the high-cost penalty of the Carbon-min solution. This supports the earlier interpretation that the balanced solution may be more practical where cost and process-hours remain important.

The breakdown of time shows the clearest construction-process mechanism. The Time-min solution reduces total process-hours by approximately 6.35 h relative to S2, mainly through lower site installation time. The Carbon-min solution reduces site installation time but increases factory production time substantially, resulting in a net increase of approximately 7.75 h. This confirms that shifting work toward a higher prefabrication setting can reduce site effort, but it may increase total process-hours if the factory production component grows more than the site component falls.

Overall, the S2 Improvement and Reference-Case Interpretation check, and component breakdown provide a more honest interpretation of the optimisation result. The final Pareto front includes a substantial subset of solutions that improve upon S2 across all three modelled objectives, but it also includes carbon-focused solutions that require higher process cost or longer process-hours. The component breakdown explains these outcomes in construction terms: improvements arise mainly through reduced setup, reduced site labour, reduced A5 installation carbon, and reduced site installation time, while penalties arise when factory processing effort and factory production time increase. This confirms that the optimisation result is not only a numerical trade-off surface, but a physically interpretable CFS framing decision-support output.

5.7.1: Practical Significance of S2-Relative Improvements

The S2-relative improvement check separates nondominated solutions that improve all three objectives from those that represent trade-offs. A solution is classified as an all-objective improvement relative to S2 only when its modelled CFS framing process cost, CFS framing A1–A5 embodied carbon, and CFS framing process-hours are all lower than the corresponding S2 reference values. This classification is deliberately stricter than simply belonging to the Pareto front.

Of the 8,302 feasible nondominated solutions, 2,695 produced lower cost, lower carbon, and lower process-hours than the S2 reference configuration. This is a strong result because it shows that S2 is not located on the best achievable cost–carbon–process-hour boundary within the implemented response model. However, the remaining Pareto solutions should not be interpreted as failures. They represent trade-off alternatives, especially carbon-priority configurations that reduce embodied carbon further but require higher cost, longer process-hours, or both.

The 2,695 all-objective improvement solutions do not mean that cost, carbon, and process-hours are non-conflicting objectives. They show that the model-calculated S2 reference is not located on the most efficient region of the implemented response surface. Trade-offs remain evident because the carbon-minimum solution requires cost and process-hour penalties, while the lowest-cost and lowest-process-hour representatives occupy different regions of the archive. The result should therefore be interpreted as an S2-relative screening finding: the framework identifies configurations that improve on the reference case and separates them from priority-specific trade-off alternatives.

5.8: TOPSIS and Stakeholder Ranking

The Pareto front provides a set of nondominated CFS framing alternatives, but the final choice depends on project priorities. A fabricator may prioritise process cost, a client or consultant may place greater emphasis on embodied carbon, and a site delivery team may prioritise reduced process-hours.

TOPSIS was used as a post-optimisation ranking layer to translate the nondominated archive into stakeholder-priority alternatives. The weighting profiles in Table 5-X were not derived from a stakeholder survey; they were defined as transparent decision scenarios to test how the preferred solution changes when process cost, A1–A5 carbon, process-hours, or balanced performance is prioritised.

Table 5-14 TOPSIS stakeholder-priority weighting profiles.

Profile	Process-cost weight	A1–A5 carbon weight	Process-hour weight	Purpose
Cost-priority	0.60	0.20	0.20	Represents cost-led CFS framing configuration screening
Carbon-priority	0.20	0.60	0.20	Represents carbon-led CFS framing configuration screening
Process-hour-priority	0.20	0.20	0.60	Represents process-hour/productivity-led CFS framing configuration screening
Balanced	0.333	0.333	0.333	Represents equal weighting of process cost, A1–A5 carbon, and process-hours

The weights in Table 5-14 make the preference assumptions explicit before the TOPSIS-selected solutions are interpreted. TOPSIS is therefore used as a decision-support ranking method, not as the optimiser and not as a source of algorithmic novelty. The NSGA-II archive remains the optimisation output; TOPSIS only ranks alternatives within that archive according to the selected stakeholder-priority profile.

Table 5-15 reports the highest-ranked Pareto solution for each stakeholder profile. The results are compared with the model-calculated S2 reference scenario, where negative percentage values indicate improvement relative to S2.

Table 5-15 TOPSIS-selected stakeholder-priority solutions.

Stakeholder profile	Cost priority	Carbon priority	Time priority	Balanced
Selected solution	P763	P1254	P2065	P1504

Stakeholder profile	Cost priority	Carbon priority	Time priority	Balanced
$(\boldsymbol{R_p})$	0.7034	0.7003	0.7105	0.7016
$(\boldsymbol{R_t})$	0.6926	0.6863	0.6839	0.7069
Panel scheme	2	2	2	2
Truss scheme	2	2	2	2
$(\boldsymbol{OC})$	0.998	1	0.811	0.898
$(\boldsymbol{P_f})$	0.7783	0.7916	0.7825	0.7817
Cost (NZD)	18,087.00	18,132.57	18,216.53	18,157.27
Carbon (kgCO_{2e})	5,271.19	5,264.52	5,289.40	5,280.55
Time (h)	88.69	88.77	87.68	88.14
Cost vs S2	−1.36%	−1.11%	−0.65%	−0.97%
Carbon vs S2	−1.78%	−1.91%	−1.44%	−1.61%
Time vs S2	−5.57%	−5.48%	−6.65%	−6.16%
TOPSIS score	0.9688	0.8945	0.9689	0.942

Note: Component breakdowns explain modelled mechanisms within the implemented response model. They are not contractor cost breakdowns, certified carbon inventories, or measured productivity records.

The TOPSIS scores are interpreted within each stakeholder-priority profile and are not used to claim that one stakeholder profile is objectively superior to another. The purpose of the ranking is to show how different decision priorities select different members of the final Pareto front.

All four TOPSIS-selected solutions improve on S2 in the three reported objectives. This does not mean they are absolute objective minima. Rather, they are compromise solutions selected from the existing nondominated front. Their value is practical: they avoid the stronger penalties associated with extreme objective choices while still improving process cost, A1–A5 embodied carbon, and process-hours relative to S2. The selected cases also show a consistent configuration pattern. All four occur in panel Scheme 2 / truss Scheme 2, with $R_p$and $R_t$located in the moderate-to-strong rationalisation band identified earlier. This indicates that the stakeholder-ranked solutions are drawn from the stable high-performing region of the Pareto front rather than from isolated extreme points.

Overall, the TOPSIS ranking demonstrates how the Pareto front can be translated into stakeholder-specific CFS framing decisions. The ranking interpreted as a transparent decision-support layer, not as a replacement for professional judgement. It provides a structured starting point for comparing CFS framing configurations under different project priorities, while retaining the full Pareto front for further design, constructability, and procurement review.

5.9: Chapter Summary

Chapter 5 showed that the reported LHS-initialised NSGA-II run produced an internally stable and interpretable nondominated archive for the implemented CFS framing response model. The convergence diagnostics, multi-seed stability results, solution-count audit, and grid-reference comparison support the use of the final archive as a model-based approximation under the stated settings.

The main result is the separation of the archive into two decision regions. Of the 8,302 nondominated solutions, 2,695 improved CFS framing process cost, CFS framing A1–A5 carbon, and CFS framing process-hours simultaneously relative to the model-calculated S2 reference configuration. The remaining 5,607 nondominated solutions were not rejected; they represent trade-off alternatives, including carbon-priority configurations that require cost or process-hour penalties.

The representative-solution analysis showed that the deterministic improvements are modest rather than transformative. The cost-minimum, process-hour-minimum, and balanced representatives support practical all-objective or near-balanced interpretation, while the carbon-minimum representative demonstrates the cost and process-hour penalty associated with minimising deterministic carbon alone. The results therefore support the use of the framework as an early-stage CFS framing screening tool rather than as a claim of universally optimal structural design.

The redesigned S2-relative classification, rationalisation-zone heatmaps, delta-from-S2 component breakdowns, and TOPSIS ranking provide the basis for the sensitivity, uncertainty, robustness, and dashboard interpretation presented in Chapters 6 and 7.

Chapter 6: Sensitivity, Uncertainty, and Robustness Results

6.1: Robustness Framework Overview

Chapter 6 evaluates the robustness and sensitivity of the implemented CFS framing response model after the final nondominated archive has been generated. The purpose is to identify which design variables and coefficient assumptions most influence CFS framing process cost, CFS framing A1–A5 carbon, and process-hours, and to test whether selected representative Pareto solutions remain interpretable under bounded uncertainty. The chapter reports OAT sensitivity, LHS-based global sensitivity, Spearman and PRCC screening, Monte Carlo robustness testing, multi-seed optimisation stability, and an OC-fixed diagnostic. These results are interpreted as model-based sensitivity and robustness evidence, not as external empirical validation.

Table 6-1 summarises the role of each robustness check. OAT sensitivity is used as a local diagnostic around the S2 reference vector. LHS-based screening and rank-based Spearman/PRCC results are used to identify broader variable-influence patterns across the bounded response space. Monte Carlo analysis tests selected representative solutions under paired coefficient uncertainty relative to S2. Multi-seed testing assesses stochastic repeatability of the NSGA-II approximation, while the OC-fixed diagnostic checks whether the reported S2-relative improvement is dependent mainly on reduced opening/detailing complexity.

Table 6-1 Robustness method roles in Chapter 6.

Robustness check	Main question answered	Main evidence used	Interpretation role
OAT sensitivity	What is the local effect of changing one variable from the S2 setting?	OAT low-bound and high-bound outputs	Local response of cost, carbon, and process-hours
LHS/rank-based screening	Which variables influence model outputs across the bounded design space?	LHS screening outputs	Broad variable-influence evidence
Spearman / PRCC screening	Which monotonic rank-based relationships are strongest?	Spearman and PRCC coefficient	Screening-level monotonic association, screening-level association
Grouped scheme comparison	How do derived panel_scheme and truss_scheme categories differ on average	Grouped mean objective shifts	Reporting-category comparison only
Monte Carlo uncertainty	Do representative Pareto solutions remain favourable under uncertain coefficients	P5/P50/P95, CV, probability relative to paired S2	Robustness of selected representative alternatives
Multi-seed stability	Is the optimisation result repeatable across independent seeds?	HV, IGD, seed-level Pareto outputs	Stochastic stability of NSGA-II approximation
OC-fixed diagnostic	Are S2-relative improvements mainly caused by low OC values	OC-fixed comparison table	Diagnostic check on opening/detailing effect

OAT sensitivity explains local behaviour, while LHS and rank-based screening provide broader variable-influence evidence. Monte Carlo analysis provides the uncertainty test for selected representative solutions. Multi-seed stability addresses a separate issue: whether the optimisation procedure gives consistent results across repeated runs. The sensitivity and Monte Carlo analyses are interpreted as robustness and uncertainty-propagation checks, not as empirical validation of commercial pricing, certified carbon performance, or measured construction productivity.

6.2: One-at-a-Time Sensitivity

One-at-a-time sensitivity was used to examine the local response of the objective functions to each design variable. In this test, one design variable was moved to its low or high bound while the remaining variables were held at the S2 reference setting. The S2 reference vector was $[R_p, R_t, OC, P_f] = [0.0, 0.0, 1.0, 0.72]$. The purpose of the test was to identify which variables have the strongest direct influence on CFS framing process cost, CFS framing A1–A5 embodied carbon, and CFS framing process-hours near the reference condition.

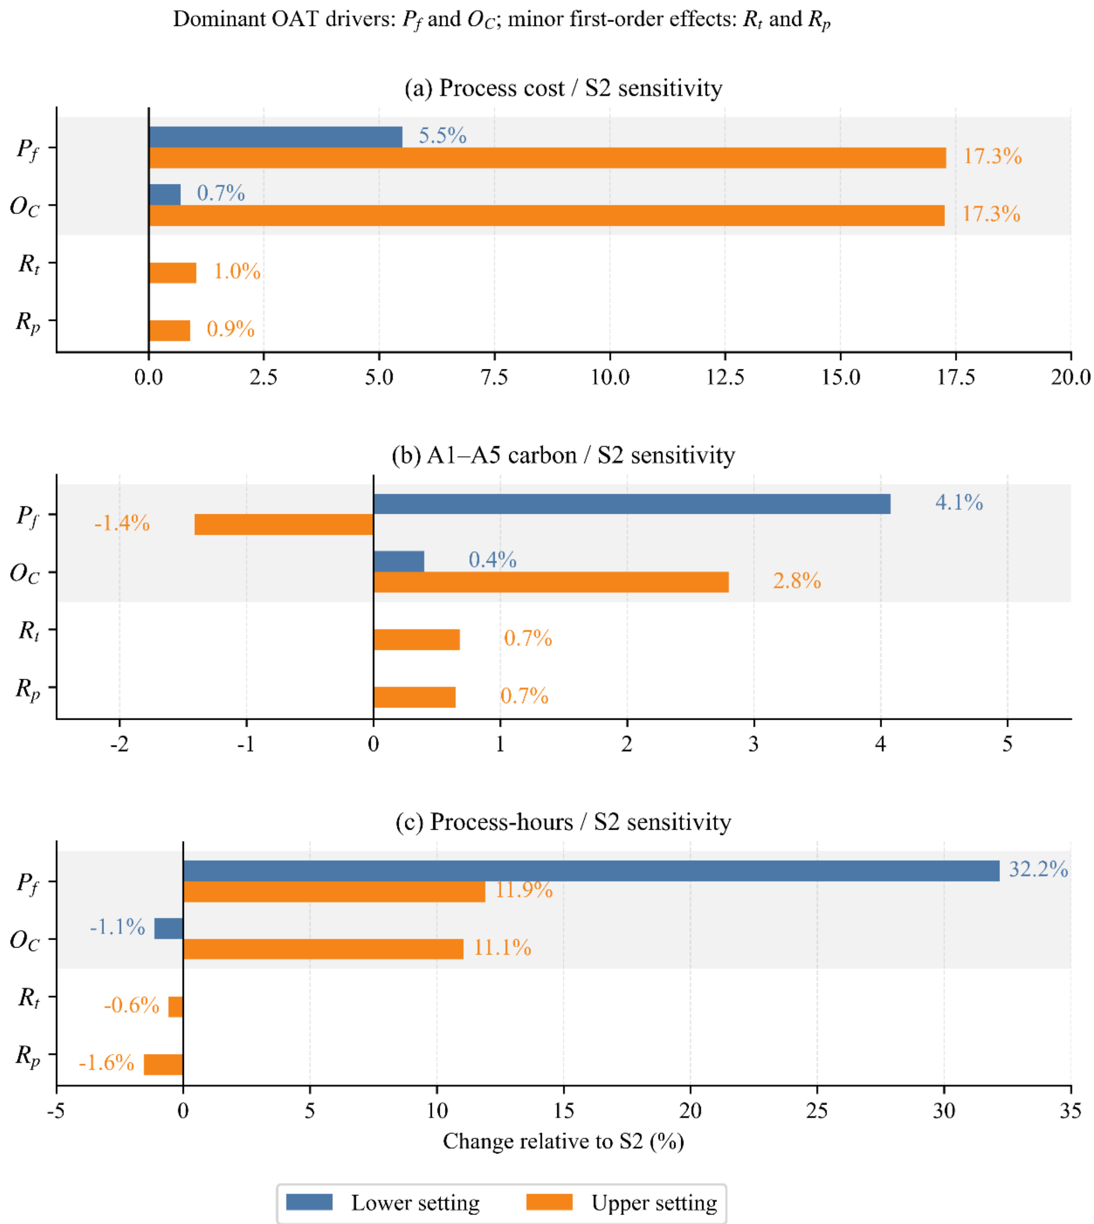

Figure 6-1 One-at-a-time sensitivity of S2-normalised objectives to decision-variable changes.

The figure shows the local response of CFS framing process cost, A1–A5 carbon, and process-hours to changes in panel rationalisation (R_p), truss rationalisation (R_t), opening/detailing complexity (OC), and prefabrication level (P_f)

Figure 6-1 shows that the implemented model response is not governed equally by all four variables. The strongest effects are associated with opening/detailing complexity and prefabrication level, while panel and truss rationalisation influence the objectives through more bounded repetition and grouping mechanisms. This supports the interpretation that the optimisation results are driven by coupled production and detailing effects rather than by a single independent variable.

Table 6-2 OAT sensitivity changes relative to the S2 reference setting.

Variable	Cost low case	Cost high case	Carbon low case	Carbon high case	Time low case	Time high case	interpretation
$(\boldsymbol{R_p})$, panel rationalisation	0.00%	+0.91%	0.00%	+0.65%	0.00%	−1.56%	Higher panel rationalisation slightly reduces process-hours, with small cost and carbon penalties
$(\boldsymbol{R_t})$, truss rationalisation	0.00%	+1.04%	0.00%	+0.68%	0.00%	−0.57%	Truss rationalisation has a weaker local time benefit than panel rationalisation
$(\boldsymbol{OC})$, opening complexity	+0.70%	+17.26%	+0.40%	+2.80%	−1.13%	+11.06%	High opening/detailing complexity strongly increases cost and process-hours
$(\boldsymbol{P_f})$, prefabrication level	+5.50%	+17.29%	+4.08%	−1.41%	+32.19%	+11.90%	Prefabrication level is the strongest local driver, especially for process-hours

Elasticity is not calculated for (R_p) and (R_t) because their S2 reference values are zero. The reported high-bound OAT responses for these two variables therefore represent full-range changes from 0.00 to 1.00 rather than local marginal derivatives at S2. These full-range changes include both the beneficial standardisation region and the nonlinear penalty region introduced to represent oversizing, trimming, handling, setup, and logistics effects. The (R_p) and (R_t) results should therefore be interpreted as net full-range rationalisation responses.

To reduce the risk of interpreting wider variable movement as stronger physical influence, the OAT results were also considered relative to the normalised movement of each variable from its S2 reference setting:

$$S_{i,j}^{range} = \frac{\Delta f_j / f_{j,S2}}{|x_i - x_{i,S2}| / (x_{i,\max} - x_{i,\min})}$$

where ($S_{i,j}^{range}$) is the range-normalised sensitivity of objective (j) to variable (i), (Δf_j) is the change in the objective value relative to S2, ($f_{j,S2}$) is the S2 objective value, (x_i) is the tested variable value, ($x_{i,S2}$) is the S2 reference value, and ($x_{i,\max} - x_{i,\min}$) is the allowable range of the variable. This index does not convert the OAT test into a global variance-decomposition method, but it reduces the risk of overstating the influence of variables simply because they move through a wider interval from S2.

The range-normalised interpretation confirms that (R_p) and (R_t) have relatively small net full-range effects compared with (OC) and (P_f). The high-(R_p) case changes cost, carbon, and process-hours by approximately (+0.91%), (+0.65%), and (-1.56%) per unit normalised movement, while the high-(R_t) case changes the same objectives by approximately (+1.04%), (+0.68%), and (-0.57%). By contrast, the high-(OC) case produces stronger penalties, with approximately (+28.77%) cost, (+4.67%) carbon, and (+18.43%) process-hour change per unit normalised movement. The prefabrication-level response is also substantial: the low- (P_f) case produces the strongest process-hour effect, while the high- (P_f) case shows the cost and time penalty associated with excessive factory/logistics burden. Overall, the OAT results indicate that, within the implemented response model and tested variable bounds, process-hours are most strongly affected by prefabrication-level movement, while cost is most strongly affected by high opening/detailing complexity after range normalisation. This supports the interpretation that (OC) and (P_f) are influential response-model variables around S2, but it should not be read as empirical proof that these variables dominate all CFS projects.

6.3: LHS-Based Design-Space Screening

The one-at-a-time sensitivity analysis in Section 6.2 examined local model behaviour near the S2 reference condition. To examine broader design-space behaviour, an independent Latin Hypercube Sampling (LHS) screening stage was also used. The LHS-based design-space screening reported in this section is separate from the LHS initialisation used to create the generation-zero NSGA-II population. Here, LHS is used as a design-of-experiments screening method to sample the bounded response space and assess broad variable influence. The purpose is not to generate a Pareto front, but to identify which design variables are most strongly associated with variation in cost, carbon, and process-hours across the sampled design domain.

The LHS screening is used here as a design-space diagnostic. It identifies rank-based associations between the four design variables and the three objective responses, without claiming variance decomposition or empirical validation.

A total of 2,000 LHS screening samples were used. This is conservative for a four-variable response model, but it was selected to provide stable rank-correlation estimates, reduce sensitivity to random sampling noise, and ensure adequate coverage of interaction regions across R_p, R_t, OC, and P_f.

Panel_scheme and truss_scheme were used only for grouped reporting because they are derived labels, whereas $R_p$and $R_t$are the optimiser-controlled rationalisation variables.

Table 6-3 Spearman rank sensitivity from LHS design-space screening.

Objective	Strongest relationship	Spearman (ρ)	Secondary relationship	Spearman (ρ)	Lower-influence variables
CFS framing process cost	(OC)	+0.757	(P_f)	+0.152	(R_t) +0.073; (R_p) −0.017
CFS framing A1–A5 carbon	(P_f)	−0.843	(OC)	+0.424	(R_p) −0.041; (R_t) +0.037
CFS framing process-hours	(P_f)	−0.590	(OC)	+0.448	(R_p) −0.122; (R_t) −0.008

Table 6-3 shows that opening/detailing complexity is the strongest rank-based driver of process cost. The positive relationship means that higher $OC$is associated with higher cost across the LHS design space. This is consistent with CFS framing practice because complex openings, non-standard lintels, custom jamb details, trimming, and coordination requirements increase both fabrication and site effort.

For embodied carbon and process-hours, P_f is the strongest Spearman-ranked variable. The negative relationship for carbon indicates that higher prefabrication level is associated with lower A1–A5 carbon across the sampled design space. The negative relationship for process-hours indicates that, across the broader design space, higher prefabrication reduces process-hours. This does not contradict the local OAT result in Section 6.2. Instead, it indicates a non-linear prefabrication response: prefabrication can reduce site work overall, but remarkably high prefabrication may still introduce factory processing and logistics penalties in some cases.

Because both metrics are rank-based, they are interpreted as screening indicators rather than direct physical sensitivity.

Table 6-4 PRCC sensitivity from LHS design-space screening.

Objective	Strongest partial relationship	PRCC	Secondary partial relationship	PRCC	Lower-influence variables
CFS framing process cost	(OC)	+0.769	(P_f)	+0.259	(R_t) +0.029; (R_p) −0.019
CFS framing A1–A5 carbon	(P_f)	−0.920	(OC)	+0.751	(R_p) −0.049; (R_t) +0.046
CFS framing process-hours	(P_f)	−0.652	(OC)	+0.544	(R_p) −0.154; (R_t) −0.052

PRCC is not applied to panel_scheme or truss_scheme as continuous sensitivity variables because these are ordinal reporting labels derived from R_p and R_t, not optimiser-controlled continuous variables. Treating them as continuous PRCC inputs would imply an artificial metric distance between scheme labels. Grouped mean comparisons are therefore used instead to describe reporting-category effects.

The PRCC results in Table 6-4 support the same overall interpretation as the Spearman screening. OC remains the dominant cost-related variable, while P_f remains the dominant variable for carbon and process-hours. The agreement between the two screening methods strengthens the interpretation that the broader model response is governed by prefabrication level and opening/detailing complexity.

Tables 6-3 and 6-4 provide the numerical Spearman and PRCC results. To make the sensitivity hierarchy clearer, Figure 6-2 presents the PRCC results as a ranked bar chart. This figure is included because the engineering interpretation depends not only on statistical values, but on whether the dominant variables can be identified quickly.

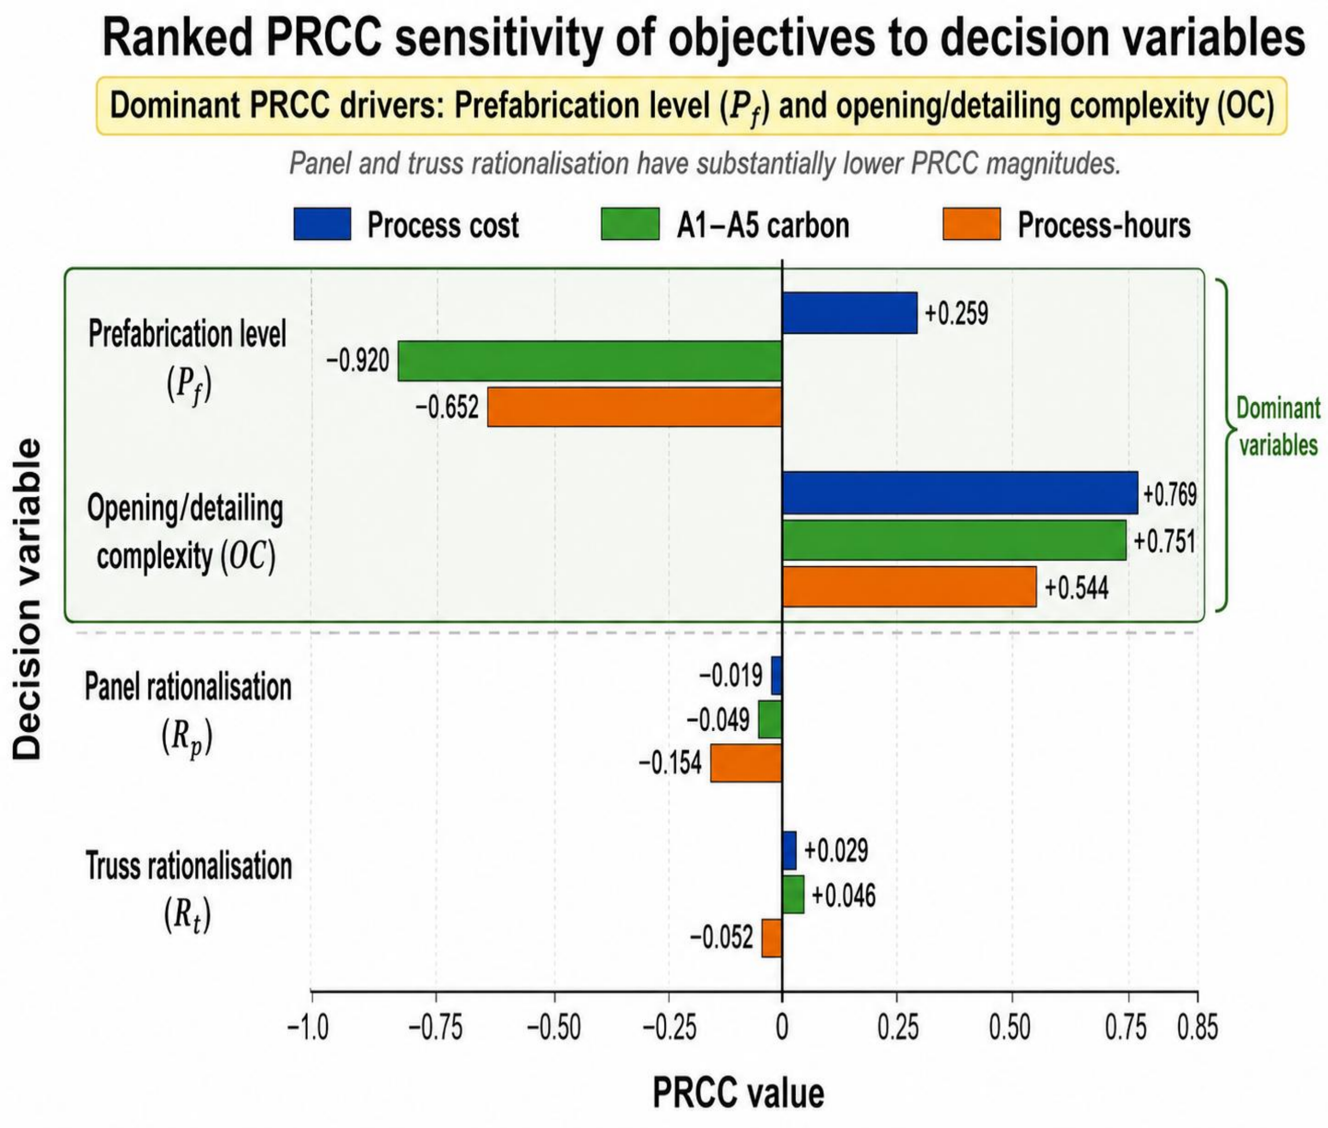

Figure 6-2 Ranked PRCC sensitivity of CFS framing objectives to decision variables.

The PRCC results support the conclusion that opening/detailing complexity and prefabrication level are the dominant variables in the implemented response model. This does not mean that panel and truss rationalisation are unimportant; rather, their effects are more bound because rationalisation benefits and penalties are both represented in response formulation.

The strong negative PRCC between P_f and carbon is interpreted as a model-specific dominance result for this case-study coefficient set, not as evidence that rationalisation has no carbon relevance in CFS framing

Overall, LHS screening supports OAT interpretation while adding broader design-space evidence. OC is the dominant cost-related variable, while P_f is the dominant carbon and process-hour variable. R_p and R_t have weaker screening effects, but they remain important for explaining standardisation branches and the structure of retained nondominated alternatives.

6.4: Monte Carlo Uncertainty Quantification

Monte Carlo analysis is used here as an uncertainty-propagation and robustness test, not as a sensitivity-ranking method or an additional optimiser. Each selected representative Pareto solution is evaluated repeatedly under sampled uncertain coefficients and compared with the S2 reference configuration under the same coefficient sample. This paired structure ensures that candidate and S2 are affected by the same uncertainty realisation, allowing the probability of S2-relative improvement to be interpreted more fairly than if the two cases were sampled independently (Helton & Davis, 2003).

The Monte Carlo analysis used $N_{MC} = 1000$ samples for each tested representative solution. A separate sample-size convergence study was not undertaken; therefore, the Monte Carlo outputs are interpreted as practical robustness estimates within the stated uncertainty ranges rather than as formally converged probabilistic results. This is considered acceptable for the decision-support purpose of the study because the analysis compares representative Pareto alternatives against paired S2 under consistent coefficient perturbations, rather than producing statistically certified industry probabilities.

The Monte Carlo results are presented as distributions because robustness depends on both central value and uncertainty spread. A representative solution may appear attractive deterministically but become less defensible if its uncertainty distribution overlaps strongly with S2 or if its probability of outperforming S2 is low. Figure 6-2 therefore compares S2 and the representative Pareto solutions using median values, P25–P75 boxes, P5–P95 whiskers, and probability of outperforming S2.

Figure 6-3 Monte Carlo distributions for S2 and representative Pareto solutions

Figure 6-3 shows that deterministic objective minima are not automatically the most robust practical choices. The carbon-minimum representative produces the lowest carbon distribution, but it also shows clear cost and process-hour penalties relative to S2. In contrast, the cost-minimum and process-hour-minimum representatives provide more balanced robustness behaviour, with stronger probabilities of outperforming S2 in the objectives most relevant to early-stage decision support. The figure should therefore be read together with Table 6-5, which reports the percentile values, coefficient of variation, and probability of outperforming S2 numerically.

Table 6-5 is the central decision table in Chapter 6 because it reports not only deterministic representative performance, but also uncertainty spread and probability of outperforming S2. This is more useful for early-stage decision support than ranking alternatives by deterministic objective values alone.

Table 6-5 Monte Carlo robustness ranking for representative solutions.

Robust rank	Solution	Cost P50 [P5–P95] NZD	Carbon P50 [P5–P95] $kgCO_{2e}$	Time P50 [P5–P95] hr	CV cost / carbon / time	Probability better than paired S2: cost / carbon / time / all three	Robust score
1	Cost-min	18,659 [16,472–20,879]	5,413 [3,993–6,844]	90.4 [88.8–92.3]	7.42% / 17.23% / 1.17%	100.0% / 100.0% / 100.0% / 100.0%	0.9400
2	Time-min	18,792 [16,495–21,171]	5,505 [3,995–6,848]	87.6 [85.4–89.9]	7.83% / 16.84% / 1.59%	69.9% / 100.0% / 100.0% / 69.9%	0.9515
3	Knee/ balanced	19,116 [16,804–21,420]	5,333 [3,958–6,808]	90.0 [87.2–92.8]	7.59% / 17.00% / 1.94%	17.1% / 100.0% / 100.0% / 17.1%	0.9861
4	Carbon-min	22,069 [18,993–25,041]	5,424 [3,963–6,751]	102.0 [97.9–105.9]	8.83% / 16.61% / 2.53%	0.0% / 100.0% / 0.0% / 0.0%	1.1207

Table 6-5 shows that the deterministic carbon-minimum solution should not automatically be treated as the best practical solution. Although it has the lowest deterministic carbon value, its cost and process-hour penalties and uncertainty behaviour make it a priority-specific alternative rather than a robust all-objective recommendation. The cost-minimum and process-hour-minimum solutions provide more defensible practical options when S2-relative robustness is considered.

The Monte Carlo results show that the absolute carbon uncertainty range is much wider than the deterministic carbon spread across the Pareto front. For example, the representative carbon P5–P95 ranges are approximately 2,700–2,900 $kgCO_2e$ wide, while the deterministic Pareto carbon spread is only 81.34 $kgCO_2e$. Therefore, the carbon objective is useful for identifying broad carbon-priority regions and S2-relative carbon improvement under paired uncertainty, but small deterministic carbon differences between nearby Pareto alternatives should not be treated as practically decisive. This strengthens, rather than weakens, the decision-support interpretation: carbon-priority decisions require robustness review rather than deterministic ranking alone

The Time-min solution is also robust, but not as strongly as the Cost-min solution. It outperforms paired S2 in carbon and process-hours in 100% of the Monte Carlo runs, but only in 69.9% of runs for cost. This means that the Time-min solution is a reliable process-hour reduction option, but its cost advantage is more sensitive to uncertain pricing and processing assumptions. From a practical CFS framing perspective, this solution is suitable where programme efficiency is important, but its cost benefit treated with moderate caution.

The Knee/balanced solution performs well for carbon and process-hours, with 100% probability of outperforming paired S2 in both objectives. However, it outperforms paired S2 in cost in only 17.1% of runs. This indicates that the balanced solution is not a cost-robust option under uncertainty, even though it remains favourable for carbon and time. It should therefore be interpreted as a compromise solution rather than a robust all-objective improvement.

The Carbon-min solution gives the weakest robustness ranking. Although it outperforms paired S2 in carbon in 100% of runs, it does not outperform paired S2 in cost or process-hours. This confirms the deterministic trade-off observed in Chapter 5: the lowest-carbon representative requires acceptance of higher process cost and longer process-hours. It is therefore a specialist carbon-priority option, not a robust within the tested uncertainty ranges across all three objectives.

The coefficient of variation values also provides useful interpretation. Carbon has the highest CV, approximately 16–17% across all representative solutions, because the carbon result is strongly affected by uncertain steel, A5 installation, waste, and freight carbon assumptions. Cost has moderate variation, with CV values of approximately 7–9%. Time has the lowest variation, with CV values below 3%, indicating that the process-hour response is more stable under the sampled uncertainty assumptions than the carbon response. Because the carbon CV is much larger than the deterministic carbon differences between several representative solutions, small carbon differences should not be over-interpreted as practically distinguishable under uncertainty. The carbon-minimum solution is therefore useful as a deterministic objective-space extreme, but not as proof of a robust carbon advantage over nearby alternatives.

Overall, the Monte Carlo analysis strengthens the interpretation of the Chapter 5 Pareto results. It shows that not all representative solutions are equally robust. The Cost-min solution is the strongest uncertainty-tested option because it remains better than S2 across all three objectives in every paired Monte Carlo run. The Time-min and Knee/balanced solutions remain dependable for carbon and process-hours but are less robust for cost. The Carbon-min solution remains a valid carbon-reduction option, but it requires clear acceptance of cost and process-hour penalties. This confirms that the decision-support value of the Pareto front lies not only in identifying low objective values, but also in showing which solutions remain defensible under uncertainty.

6.5: Multi-Seed Stability Check

The Monte Carlo analysis in Section 6.4 tested uncertainty in model assumptions such as cost rates, carbon factors, productivity, and prefabrication response. A separate multi-seed stability check was used to test the repeatability of the optimisation result itself. This distinction is important because variation in construction assumptions is different from variation caused by the random starting conditions of the optimisation process.

The multi-seed check repeated the same optimisation configuration using five independent seeds: 11, 22, 33, 44, and 55. For each run, the final trade-off coverage, distance from the combined reference front, number of nondominated solutions, and best objective values were recorded. The purpose was to confirm whether the same cost–carbon–process-hour region could be recovered consistently, rather than relying on one favourable run.

The multi-seed evidence is summarised here only to connect optimisation repeatability with Chapter 6 robustness interpretation.

Table 6-6 Multi-seed stability summary.

Metric	Mean	Standard deviation	Minimum	Maximum	CV (%)
Final hypervolume	0.04233028	0.00000015	0.04233007	0.04233043	0.00035
IGD against multi-seed reference	0.00004612	0.00000313	0.00004290	0.00005065	6.78370
Final nondominated solutions	8,201.40	466.31	7,392.00	8,502.00	5.68576
Minimum normalised cost	0.98391274	0.00000019	0.98391260	0.98391307	0.00002
Minimum normalised carbon	0.97409071	0.00000060	0.97409033	0.97409176	0.00006
Minimum normalised process-hours	0.93241225	0.00003403	0.93239064	0.93247238	0.00365

The number of nondominated solutions varied from 7,392 to 8,502. This variation does not indicate a weak result. In this type of optimisation, different runs can produce different densities of solutions along the same trade-off region. The more important finding is that the objective limits and overall trade-off coverage remained stable.

The IGD result also supports this interpretation. The mean IGD against the combined multi-seed reference front was 0.00004612, showing that each seed produced a front close to the combined reference result. The higher coefficient of variation for IGD is expected because this measure is more sensitive to local spacing between points, while hypervolume reflects the overall coverage of the trade-off space.

Overall, the multi-seed check supports the repeatability of the reported NSGA-II approximation. Although the number of nondominated solutions varied between seeds, the objective limits, hypervolume, IGD, and trade-off coverage remained stable. This strengthens confidence that the sensitivity and robustness interpretation is not dependent on one favourable optimisation seed.

6.6: OC-Fixed Diagnostic

The OC-fixed diagnostic was added to test whether the Chapter 5 improvement pattern was mainly caused by allowing opening/detailing complexity to reduce below the S2 reference condition. This is an important diagnostic because $OC$is one of the strongest sensitivity variables and could otherwise make the optimisation appear to improve all objectives simply by simplifying detailing.

In this diagnostic, selected Pareto solutions were re-evaluated with OCfixed at the S2 reference value while retaining their panel rationalisation, truss rationalisation, and prefabrication-level settings. This isolates the contribution of rationalisation and prefabrication from the benefit associated with simplified opening/detailing conditions.

This check is important for CFS framing interpretation. Opening/detailing complexity represents the standardisation of lintels, jamb studs, trimming, service penetrations, and installation detailing. If the Pareto improvement disappeared when OC was fixed at 1.0, then the results would depend heavily on opening/detailing simplification. If the improvement remained, then the results would be supported by broader configuration effects, including rationalisation and prefabrication level.

Table 6-7 OC-fixed diagnostic results.

Diagnostic measure	Original Pareto front	OC fixed at 1.0	Interpretation
Total evaluated Pareto solutions	8,302	8,302	Same Chapter 5 Pareto solutions were re-evaluated
Solutions with cost below or equal to S2	2,695	2,694	Almost unchanged
Solutions with carbon below or equal to S2	8,302	8,302	Carbon improvement does not depend on low (OC)
Solutions with process-hours below or equal to S2	5,319	5,318	Almost unchanged
Solutions better than S2 in all three objectives	2,695	2,694	Simultaneous improvement remains after fixing (OC)
Cost range relative to S2	−1.61% to +15.58%	−1.61% to +15.58%	Cost range is unchanged
Carbon range relative to S2	−2.59% to −1.08%	−2.59% to −1.08%	Carbon range is unchanged
Process-hour range relative to S2	−6.76% to +8.25%	−5.63% to +8.25%	Maximum time saving reduces, but remains better than S2

Table 6-7 shows that fixing OC at the S2 reference value has only a minor effect on the overall improvement pattern. The number of solutions better than S2 in all three objectives changes from 2,695 to 2,694. Carbon improvement is unchanged across the full nondominated set, and the number of process-hour-improving solutions changes by only one solution.

The OC-fixed results show that the improvement pattern is not explained solely by reducing opening/detailing complexity. Holding $OC = 1.0$removes one major simplification pathway, but the model still identifies improvement and trade-off behaviour through rationalisation and prefabrication effects. This supports the interpretation that the Chapter 5 archive is driven by coupled configuration mechanisms rather than by a single low-OC assumption.

The main effect of fixing OC is on the best process-hour result. The maximum process-hour reduction changes from −6.76% to −5.63% relative to S2. This indicates that low opening/detailing complexity contributes to the strongest time-saving case, but it is not the main reason the Pareto front performs better than S2. The broader improvement remains supported by the combined effects of rationalisation and prefabrication level.

6.7: Robustness Summary

Chapter 6 showed that the implemented response model is most sensitive to opening/detailing complexity and prefabrication level, with panel and truss rationalisation producing more bounded effects through repetition and grouping mechanisms. The addition of a PRCC ranked bar chart makes this hierarchy visible beyond the numerical tables.

The Monte Carlo results showed that uncertainty is not evenly distributed across objectives. Carbon had the largest uncertainty spread, cost showed moderate spread, and process-hours were comparatively stable. This supports the interpretation that carbon-priority decisions should be treated cautiously when based on generic or uncertain emission factors.

The robustness ranking showed that deterministic objective minima are not automatically the most practical choices. In particular, the carbon-minimum representative is best interpreted as a priority-specific trade-off rather than a robust all-objective recommendation. The cost-minimum and process-hour-minimum representatives provide stronger practical decision-support value when uncertainty and probability of outperforming S2 are considered.

The OC-fixed diagnostic confirmed that the Chapter 5 improvement pattern is not explained solely by reducing opening/detailing complexity. Holding OC at the S2 reference condition still allowed rationalisation and prefabrication effects to generate interpretable improvement and trade-off behaviour. Chapter 7 uses these findings to frame the dashboard as a practitioner-facing interpretation tool rather than a deterministic design selector.

Chapter 7: Discussion and Practical Decision Support

7.1: Discussion Overview

Chapter 7 discusses how the optimisation, sensitivity, uncertainty, and robustness results can be translated into practical early-stage CFS framing decision support. The purpose is not to repeat the numerical results from Chapters 5 and 6, and not to describe the dashboard interface in detail. Instead, the chapter interprets what the results mean for CFS fabricators, designers, and project decision-makers who need to compare configuration alternatives before detailed structural verification, final pricing, certified carbon assessment, and construction planning.

The main decision-support contribution is the workflow that connects project-derived CFS variables, S2-relative classification, representative-solution interpretation, TOPSIS stakeholder-priority ranking, and Monte Carlo robustness evidence. The dashboard is retained as a visual interpretation layer within this workflow, not as the primary contribution and not as an automatic design selector.

The Scenario Assumptions / Recalibration Check page provides bounded stress-testing of selected cost and carbon assumptions. It allows users to examine how selected-solution and S2-comparison values respond to changes in assumptions such as steel material cost, factory processing cost, site labour rate, transport distance, freight carbon factor, A1–A3 steel carbon factor, and A5 installation carbon factor. These controls support interpretation of the exported results under changed assumptions, but they do not rerun NSGA-II, regenerate the nondominated archive, replace the imported Pareto data, or create a new official optimisation result.

Table 7-1 Dashboard boundary table.

Dashboard aspect	Included in dashboard	Not included / not inferred
Decision variables	Interactive exploration of (R_p), (R_t), (OC), and (P_f)	Structural member sizing or redesign
Objectives	CFS framing process cost, CFS framing A1–A5 carbon, and process-hours	Whole-building cost, whole-building LCA, or calendar-duration scheduling
Pareto results	Exploration of imported nondominated alternatives and representative solutions	Evidence of one universally optimal CFS configuration
S2 comparison	Model-calculated comparison against the S2 reference scenario	Measured contractor benchmark or externally certified project baseline
Scenario controls	Recalculate selected-solution and S2 comparison values under changed cost/carbon assumptions	Does not regenerate the nondominated archive or replace the imported Python optimisation results
Scenario assumptions	Bounded stress-testing of selected cost and carbon assumptions	Full commercial recalibration or supplier quotation system
Robustness evidence	Imported sensitivity and Monte Carlo robustness outputs	Real-world performance guarantee
Decision ranking	Stakeholder-weighted ranking of existing alternatives	Automatic final design selection without professional judgement
Engineering boundary	Early-stage CFS framing configuration decision support	Consent documentation, AS/NZS 4600 verification, quantity surveying, or structural certification

The dashboard changes the interpretation of the optimisation results by turning the Pareto front into a practical decision workflow. A user can compare whether a configuration is more suitable for cost reduction, carbon reduction, process-hour reduction, or robustness under uncertainty. This is important because the research results do not identify one universal best design. Instead, they identify a set of feasible alternatives, each with different trade-offs.

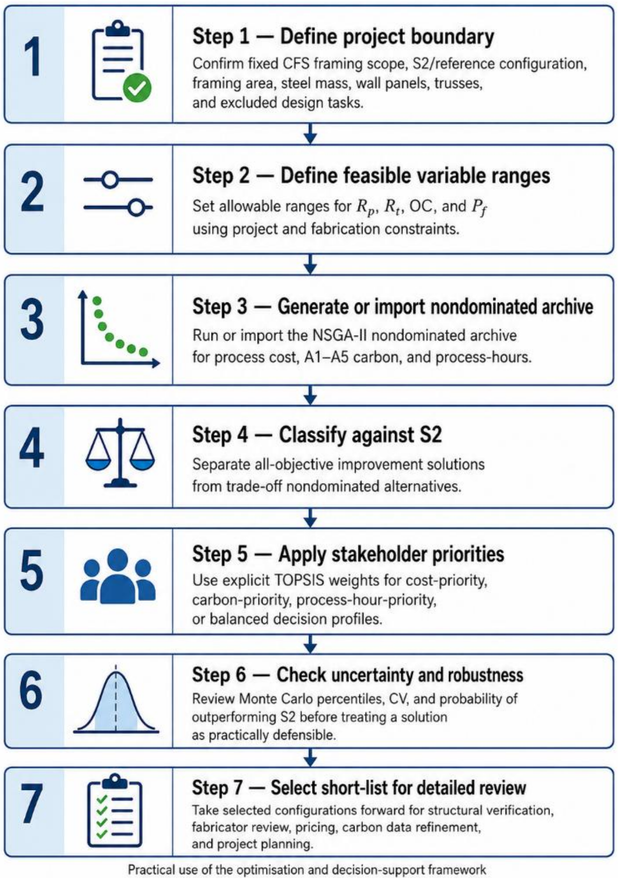

Figure 7-1 Practitioner decision workflow for early-stage CFS framing screening.

Figure 7-1 clarifies the role of the dashboard within the decision-support framework. The dashboard is not the decision-maker, and it is not the optimiser. Its role is to help users move through the workflow by viewing the archive, filtering configurations, comparing alternatives against S2, checking representative solutions, reviewing uncertainty outputs, and applying explicit stakeholder-priority weights.

The dashboard is intended for the early framing-configuration stage, after a preliminary CFS framing layout and S2-type reference quantities are available, but before final fabrication detailing, supplier tendering, and construction planning are fixed. At this stage, designers, CFS fabricators, project engineers, and sustainability assessors can still compare rationalisation, detailing, and prefabrication strategies without changing the structural system. The dashboard therefore supports decisions such as whether to prioritise lowest process cost, lowest modelled A1–A5 carbon, shortest process-hours, or a balanced/robust configuration. It does not make the final engineering decision; it narrows and explains the set of defensible options before formal design checking and procurement.

7.2: Practitioner Decision Workflow

The dashboard is retained in the discussion because it provides a practical way to inspect the optimisation and robustness outputs. However, the dashboard should be interpreted as a communication layer rather than as a separate analytical contribution. Its value is that it places Pareto alternatives, S2-relative classification, representative solutions, sensitivity results, Monte Carlo robustness, and TOPSIS ranking in one decision-support environment.

The most important dashboard function is not visual appearance, but decision traceability. A practitioner can begin with the full nondominated archive, filter alternatives by objective performance or configuration variables, compare selected solutions against S2, and then check whether the preferred solution remains robust under uncertainty. This supports a disciplined workflow in which deterministic optimisation results are not accepted until they have been tested against sensitivity and Monte Carlo evidence.

The dashboard also helps prevent misinterpretation of Pareto results. A nondominated solution is not automatically a practical improvement over the reference case, and a deterministic objective minimum is not automatically robust. By showing S2-relative classification, representative solutions, robustness outputs, and stakeholder-priority ranking together, the dashboard supports interpretation rather than automatic selection.

7.3: Practitioner Use Case for CFS Framing Configuration Screening

An illustrative practitioner use case begins with a CFS fabricator, project engineer, or sustainability assessor reviewing the imported nondominated alternatives in the Pareto Explorer. This use case is not presented as evidence of industry adoption or user-tested usability; it demonstrates how the dashboard could support early-stage interpretation of the model outputs. The user first identifies whether the project priority is lowest process cost, lowest CFS framing A1–A5 carbon, lowest process-hours, balanced performance, or robustness under uncertainty. This step is necessary because the optimisation result does not provide one universal best configuration; it provides a set of feasible trade-off alternatives.

The user then selects a representative solution or filters the nondominated archive using objective thresholds and design-variable ranges. For example, a cost-priority user may begin with the cost-minimum representative, while a contractor concerned with site labour may inspect the time-minimum representative. A sustainability-focused user may inspect the carbon-minimum representative, but the dashboard also shows whether that carbon benefit is accompanied by cost or process-hour penalties. The selected configuration is then compared against S2 using objective cards, percentage-change values, and component breakdowns.

The next step is robustness review. The user checks whether the selected solution is supported by the Monte Carlo outputs and whether its performance remains favourable relative to paired S2 under the tested coefficient ranges. This is important because a deterministic low-carbon or balanced solution may not be the most defensible option once coefficient uncertainty is considered. The dashboard therefore supports a decision sequence of filtering, comparison, robustness checking, and ranking rather than a single-step selection.

A practical use case would proceed as follows. First, the practitioner identifies the project reference configuration and confirms that the fixed boundary is comparable to S2 in scope. Second, the practitioner defines realistic ranges for panel rationalisation, truss rationalisation, opening/detailing complexity, and prefabrication level. Third, the practitioner reviews the nondominated archive and separates all-objective improvement solutions from trade-off alternatives. Fourth, the practitioner applies explicit stakeholder weights, for example cost-priority, carbon-priority, process-hour-priority, or balanced weights. Fifth, the practitioner checks Monte Carlo robustness before short-listing a configuration for detailed engineering and commercial review.

Finally, the user can use the stakeholder-weighted ranking tab to test how the preferred alternative changes when cost, carbon, process-hours, and robustness are weighted differently. This does not automate the final engineering decision. Instead, it makes the consequences of stakeholder priorities explicit and provides a traceable basis for discussing CFS framing configuration alternatives before final detailing, procurement, and construction planning.

7.4: Engineering Interpretation of Derived Panel and Truss Reporting Categories

Sections 7.4 and 7.5 interpret the engineering meaning of the optimisation findings rather than repeating the numerical results already reported in Chapters 5 and 6.

The final nondominated archive was interpreted using derived panel and truss reporting categories. These labels are decoded from the continuous rationalisation-intensity variables $R_p$and R_t. They are useful because they translate continuous model responses into practical CFS framing branch descriptions, but they are not independent optimiser-controlled variables. The engineering interpretation should therefore focus on what the retained reporting categories indicate about rationalisation intensity, repetition, fabrication standardisation, and logistics implications within the case-study model.

Within the reported case-study archive, the retained nondominated solutions are concentrated in a moderate-to-high rationalisation branch. This should be interpreted as a model-specific pattern rather than a universal fabrication rule. Most retained nondominated solutions were located in the panel Scheme 2 / truss Scheme 2 branch, while only a small number of solutions entered the panel Scheme 3 / truss Scheme 2 branch. This shows that the main trade-off surface is governed by a stable rationalisation region rather than by all possible panel and truss reporting categories. The result suggests that rationalisation can be beneficial when repetition gains are balanced against oversizing, handling, setup, and logistics effects, but it does not prove that the same branch would dominate for other building layouts, fabricator libraries, transport constraints, or structural systems.

This result does not mean that panel or truss rationalisation is unimportant. Instead, it indicates that the model identifies a narrow-preferred rationalisation range for this fixed case-study geometry. Panel rationalisation remains a useful standardisation lever, particularly for process-hour reduction and branch selection. Truss rationalisation is more stable and secondary in this case study, with the final Pareto front retaining one preferred truss reporting category while still varying R_t within that branch.

The truss Scheme 2 concentration should therefore be interpreted as a case-study-preferred truss standardisation branch, not as evidence that truss rationalisation has no practical relevance. For a fixed residential roof geometry, truss-family variation is constrained by roof layout, span arrangement, fabrication limits, and handling requirements. Once a practical truss family is selected, the remaining objective differences are more strongly influenced by opening/detailing complexity, prefabrication level, and panel rationalisation.

Engineering interpretation is that rationalisation is useful only when it improves repetition without creating excessive grouping, trimming, handling, or coordination penalties. The retained nondominated solutions indicate a preference for moderate-to-strong rationalisation rather than minimum or maximum rationalisation. This supports a practical recommendation: fabricators should rationalise panel and truss families enough to improve repetition and setup efficiency but not assume that maximum standardisation is automatically optimal.

7.5: Design-Variable Interpretation for Practice

The results indicate that OC and P_f should be treated as practical design-management variables rather than abstract optimiser settings. Opening/detailing complexity affects the burden of lintels, jambs, trimming, local fit-up, and site adjustment. Prefabrication level affects the factory–site work split, including factory preparation, handling, logistics, and site installation. Their influence in the sensitivity results is therefore consistent with the practical role of detailing and prefabrication decisions in CFS framing delivery.

For practice, the implication is not simply to minimise OC or maximise P_f. Reducing opening/detailing complexity may improve all three objectives, but only where it is compatible with architectural requirements and structural detailing. Increasing prefabrication may reduce site process-hours, but it can introduce factory, logistics, and coordination penalties. The useful decision is therefore a controlled combination of rationalisation, detailing simplification, and prefabrication level rather than an extreme setting of one variable.

This interpretation is consistent with the OC-fixed diagnostic in Chapter 6. Holding OC at the S2 reference level reduces one improvement pathway, but it does not remove the broader rationalisation and prefabrication trade-off structure. The framework therefore supports configuration screening rather than single-variable optimisation.

Table 7-2 Engineering mechanism summary for OC and P_f.

Variable / mechanism	Engineering meaning	Practical implication

Higher ($\boldsymbol{OC}$)	More opening/detailing complexity increases trimming, lintel coordination, junction detailing, and site adjustment effort	Keep openings and details as standardised as practical
Lower ($\boldsymbol{OC}$)	Represents simpler, more repeatable detailing, not removal of openings	Use repeatable opening modules and consistent junction details
Increasing ($\boldsymbol{P_f}$)	Shifts work from site to factory and can reduce site installation burden and A5 impacts	Useful for reducing process-hours and site-related carbon
Very high ($\boldsymbol{P_f}$)	Can increase factory processing, handling, logistics, and saturation effects	Maximum prefabrication should not be assumed to be best
Preferred region	Cost minimum occurs near (P_f=0.71), while time minimum occurs near (P_f=0.78)	Cost-focused and time-focused decisions may require different prefabrication levels

The practical conclusion is not that one prefabrication level is universally optimal, but that cost-focused and time-focused decisions may require different factory–site balances. In this case-study model, a moderate prefabrication level is more favourable for cost, while a slightly higher prefabrication level is more favourable for process-hours.

7.6: Practical Meaning of S2-Relative Screening

The practical meaning of S2-relative screening is that it gives the decision-maker a reference-based way to interpret a large nondominated archive. In practice, a fabricator or designer does not only need to know whether a solution is Pareto-optimal under the model. They need to know whether it is better, worse, or trade-off dependent relative to a familiar project reference.

The S2-relative classification therefore converts the archive into decision categories. All-objective improvement solutions are useful for identifying configurations that merit further review because they improve process cost, A1–A5 carbon, and process-hours relative to S2 under the implemented model. Trade-off nondominated solutions remain useful where a stakeholder is willing to accept a penalty in one objective to prioritise another. This is particularly important for carbon-priority decisions, where the deterministic carbon-minimum solution may not be the most robust practical option.

The discussion contribution is therefore not the count itself, which was reported in Chapter 5, but the decision rule: Pareto membership should be interpreted together with S2-relative performance, stakeholder priority, and robustness evidence before a configuration is short-listed.

Table 7-3 Practical decision rules derived from the optimisation and robustness results.

Decision question	Evidence to check	Practical rule
Is the solution genuinely better than S2?	S2-relative classification and objective ratios	Do not treat Pareto membership alone as improvement
Is the solution priority-specific or generally useful?	Representative solution type and component changes	Carbon-minimum and other extremes require trade-off review
Is the result robust under uncertainty?	Monte Carlo P5/P50/P95, CV, and probability better than S2	Do not short-list deterministic minima without robustness evidence
Is the solution driven only by reduced (OC)?	OC-fixed diagnostic	Check whether improvement remains when ($OC = 1.0$)

Decision question	Evidence to check	Practical rule
Is prefabrication being over-assumed?	(P_f), logistics, factory, and site-hour components	Avoid assuming maximum prefabrication is always best
Can the solution move to implementation?	Structural, pricing, carbon, and construction review	Use the framework to short-list, not to finalise design

7.7: Limitations and Practical Applicability

Chapter 7 translated the optimisation, sensitivity, uncertainty, and robustness results into practical decision-support guidance. The discussion showed that the main engineering value of the framework is not the use of NSGA-II or the dashboard interface itself, but the structured workflow connecting project-derived CFS variables, S2-relative classification, representative-solution interpretation, Monte Carlo robustness, and stakeholder-priority ranking.

The framework should be used as an early-stage screening method. It helps compare configuration alternatives before detailed structural verification, project-specific pricing, product-specific carbon assessment, and construction planning. The results are therefore most useful for narrowing the decision space and identifying promising alternatives for further review, rather than for approving a final CFS framing design.

The practitioner workflow clarifies how a CFS fabricator or designer could use the framework: define the project boundary, set realistic ranges for R_p, R_t, OC, and P_f, classify the nondominated archive relative to S2, apply explicit stakeholder weights, check robustness, and then short-list configurations for detailed review. This makes the framework a screening and decision-support process rather than an automatic design selector.

The chapter also clarified the practical meaning of the main findings. Moderate rationalisation is preferable to maximum rationalisation in the implemented model; opening/detailing simplification is important but not the only improvement pathway; prefabrication level must be balanced against factory, logistics, and site effects; and deterministic carbon-minimum solutions require robustness review before being treated as practical recommendations. Chapter 8 draws these findings into the final conclusions and recommendation.

Chapter 8: Conclusions and Recommendations

8.1: Summary of Research

This research developed and evaluated a bounded multi-objective decision-support framework for early-stage prefabricated CFS residential framing configuration. The framework compared alternative CFS framing configurations using three simultaneous objectives: CFS framing process cost, CFS framing A1–A5 carbon, and construction process-hours. The implemented decision vector was $x = [R_p, R_t, OC, P_f]$, where R_p and R_t represent continuous rationalisation-intensity variables, OC represents opening/detailing complexity, and P_f represents prefabrication level. The model was applied to a fixed residential CFS case-study boundary using a project-derived, model-calculated S2 reference configuration. The conclusions are therefore interpreted as model-based findings for the implemented case-study framework, not as universal CFS design rules, contractor pricing results, certified structural design, or whole-building LCA outcomes.

The reported LHS-initialised NSGA-II run produced a final nondominated archive of 8,302 feasible solutions with no hard constraint violations. Within this archive, 2,695 solutions improved on S2 across all three modelled objectives. The optimisation result was supported by hypervolume convergence, grid-reference comparison, solution-count audit, and multi-seed stability checks. Sensitivity analysis, LHS-based screening, Monte Carlo uncertainty testing, and the OC-fixed diagnostic further supported the interpretation that opening/detailing complexity and prefabrication level were the dominant response drivers in the implemented model, while panel and truss rationalisation had more branch-specific effects.

The dashboard translated the optimisation, sensitivity, uncertainty, TOPSIS ranking, and S2-comparison outputs into a practical decision-support workflow. Overall, the research demonstrates that CFS framing configuration can be assessed through a transparent and uncertainty-aware optimisation process, provided that the numerical findings are interpreted within the S2 case-study boundary and selected coefficient assumptions.

8.2: Answers to Research Questions

The research questions are answered below using the results reported in Chapters 5, 6, and 7. Each answer is limited to the fixed CFS framing case-study boundary, the implemented response model, and the model-calculated S2 reference configuration.

8.2.1: Answer to Research Question 1

RQ1: How do panel rationalisation intensity, truss rationalisation intensity, opening/detailing complexity, and prefabrication level affect CFS framing process cost, A1–A5 embodied carbon, and process-hours within the fixed residential case-study boundary?

The results show that CFS framing performance is affected by the combined action of rationalisation, opening/detailing complexity, and prefabrication level. The strongest conclusion is not that one variable should be maximised or minimised, but that configuration decisions create coupled factory–site, material–labour, and cost–carbon–process-hour responses. Moderate rationalisation was more favourable than extreme rationalisation because repetition benefits were retained without incurring the full grouping, oversizing, or coordination penalties represented in the response model.

The answer to RQ1 is therefore that early-stage CFS framing performance should be evaluated as a configuration problem rather than as an isolated material, cost, or prefabrication problem. R_p, R_t, OC, and P_f are useful because they allow panel repetition, truss repetition, detailing burden, and factory–site work split to be compared within one response framework.

8.2.2: Answer to Research Question 2

RQ2: What Pareto-optimal cost–carbon–time configurations are achievable for the CFS residential framing case study relative to the S2 reference configuration?

The optimisation showed that Pareto membership and practical improvement relative to S2 are not the same. The final nondominated archive contained both all-objective improvement solutions and priority-specific trade-off alternatives. This distinction is a major conclusion because it prevents the common optimisation error of treating every nondominated solution as automatically better than the reference case.

The presence of 2,695 all-objective improvement solutions does not mean that cost, carbon, and process-hours are non-conflicting. It means that the model-calculated S2 reference was not located on the most efficient region of the implemented response surface. Trade-offs remain evident because the carbon-minimum representative achieved lower deterministic carbon only with cost and process-hour penalties, while the cost-minimum, process-hour-minimum, and balanced representatives occupied different regions of the archive.

The answer to RQ2 is therefore that the dominant decision value of the Pareto archive is not simply the production of many nondominated points. Its value is the separation of practical S2-improvement regions from trade-off regions, allowing users to distinguish generally favourable configurations from priority-specific alternatives.

8.2.3: Answer to Research Question 3

RQ3: How robust are selected Pareto configurations relative to S2 under uncertainty in cost rates, carbon factors, productivity assumptions, logistics, and waste allowances?

The results show that rationalisation and prefabrication are beneficial only when controlled within the wider configuration response. Moderate panel and truss rationalisation produced favourable objective behaviour because it captured repetition and standardisation benefits without pushing the model into stronger oversizing, grouping, or coordination penalties. Similarly, prefabrication level affected process-hours through site-labour reduction, but it also introduced factory, logistics, handling, and coordination effects.

The answer to RQ3 is therefore that CFS fabricators and designers should not treat rationalisation or prefabrication as one-directional improvement levers. The practical conclusion is to seek a coordinated configuration: enough rationalisation to improve repetition, enough prefabrication to reduce site effort, and enough detailing control to avoid excessive fabrication or fit-up burden.

8.2.4: Answer to Research Question 4

RQ4: How can the optimisation and robustness results be translated into a practical dashboard-based decision-support workflow for early-stage CFS framing configuration decisions?

The uncertainty analysis showed that deterministic objective minima are not automatically the most defensible practical choices. Carbon had the largest uncertainty spread, while process-hours were comparatively stable. This means that carbon-priority decisions are more exposed to coefficient uncertainty, especially where generic steel emission factors or construction-stage carbon assumptions are used.

The Monte Carlo results also showed why robustness must be considered before short-listing a representative solution. The deterministic carbon-minimum solution was useful as a carbon-priority alternative, but its cost and process-hour penalties made it less suitable as a general practical recommendation. In contrast, the cost-minimum and process-hour-minimum representatives provided stronger practical decision-support value because their S2-relative performance was more robust across the uncertainty testing.

The answer to RQ4 is therefore that early-stage CFS framing decisions should not be based only on deterministic objective values. Robustness evidence, including percentile spread, coefficient of variation, and probability of outperforming S2, should be checked before a configuration is carried forward.

8.2.5: Overall Answer to the Research Problem

The central research problem is answered by showing that early-stage CFS residential framing configuration decisions can be evaluated more transparently when process cost, CFS framing A1–A5 carbon, and process-hours are modelled together. The results show that preferred configurations depend on the interaction between rationalisation intensity, opening/detailing complexity, prefabrication level, and coefficient assumptions. The framework does not produce one universal optimum; it produces a traceable nondominated decision set that can be filtered, compared, stress-tested, ranked, and communicated.

The overall contribution is therefore a bounded decision-support workflow for CFS residential framing configuration. Within the fixed case-study boundary, the workflow links a transparent response model, LHS-initialised NSGA-II optimisation, sensitivity screening, Monte Carlo robustness testing, representative-solution interpretation, TOPSIS ranking, and dashboard-based communication. Its value lies in making trade-offs explicit before detailed structural verification, supplier pricing, procurement, and construction planning.

8.3: Overall Engineering Conclusion

The overall engineering conclusion is that the main value of the framework lies in making configuration trade-offs visible before detailed project commitment. In the implemented case study, improved performance was not achieved by simply maximising prefabrication, minimising carbon, or selecting the most rationalised framing layout. Instead, the favourable solutions emerged from a controlled balance of rationalisation, detailing simplification, and prefabrication level.

This conclusion is important for CFS residential framing because the practical decision is rarely a pure optimisation problem. Fabricators and designers must balance material use, factory setup, site labour, logistics, detailing burden, carbon assumptions, and project-specific constraints. The thesis shows that these effects can be represented in a bounded response model and interpreted through S2-relative comparison, uncertainty testing, and stakeholder-priority ranking.

The framework should therefore be understood as a preliminary configuration-screening method. Its role is to identify promising CFS framing alternatives and expose trade-off mechanisms before final structural verification, pricing, product-specific carbon assessment, and construction planning. The results indicate that the principal trade-off exposed by the framework is not generic cost-versus-carbon opposition, but the balance between carbon reduction, prefabrication level P_f, and the factory–logistics–site process penalties that occur at high prefabrication intensity.

8.4: Original Contributions

The original contribution of this study is the structured integration of established optimisation, uncertainty, ranking, and dashboard methods into a CFS-specific decision-support framework. The original contribution lies in the CFS-specific integration of these established methods.

The main contributions are as follows.

Contribution 1: CFS-specific configuration-screening formulation.

The study formulated early-stage prefabricated CFS framing as a configuration-screening problem using R_p, R_t, OC, and P_f. This contribution matters because it shifts the decision problem away from abstract optimisation and toward the production variables that fabricators and designers can influence before detailed design finalisation.

Contribution 2: Coefficient-audited cost–carbon–process-hour response model.

The study developed a transparent response model linking configuration variables to CFS framing process cost, A1–A5 carbon, and process-hours. The coefficient audit is a key contribution because it makes the modelling assumptions visible and testable rather than hiding them inside a black-box optimiser.

Contribution 3: S2-relative interpretation of nondominated solutions.

The study separated Pareto membership from improvement relative to the model-calculated S2 reference. This contribution matters because it prevents nondominated alternatives from being misinterpreted as automatically better than the reference case.

Contribution 4: Robustness-aware representative-solution interpretation.

The study combined representative Pareto solutions with Monte Carlo robustness testing. This contribution matters because it shows that deterministic objective minima are not always the most defensible practical recommendations under uncertainty.

Contribution 5: Practitioner-facing decision-support workflow.

The study translated optimisation and robustness outputs into a dashboard-supported decision workflow. This contribution matters because it gives CFS practitioners a structured way to move from archive exploration to weighted ranking, robustness checking, and short-listing for detailed review.

8.5: Recommendations for Practice

Recommendation 1: Use S2-relative classification before interpreting Pareto solutions. Practitioners should not treat nondominated solutions as automatically better than the reference case. Chapter 5 showed that the archive contained both all-objective improvement solutions and trade-off alternatives. Therefore, CFS fabricators and designers should first classify candidate solutions relative to a project reference configuration before short-listing them.

Recommendation 2: Target controlled rationalisation rather than maximum rationalisation.

The heatmap results showed that the favourable response region occurred around moderate panel and truss rationalisation, not at the maximum rationalisation boundary. Practitioners should therefore use rationalisation to improve repetition and standardisation but avoid assuming that maximum grouping or standardisation is always beneficial.

Recommendation 3: Treat opening/detailing complexity as a design-management variable.

Sensitivity and OC-fixed diagnostic results showed that OC strongly affects the response but does not explain all improvement alone. Practitioners should review openings, lintels, jambs, trimmers, local reinforcement, and fit-up burden early because detailing simplification can influence cost, carbon, and process-hours.

Recommendation 4: Do not maximise prefabrication without checking factory–site trade-offs.

The prefabrication results showed that higher P_f can reduce site process-hours but can also introduce factory, logistics, handling, and coordination penalties. Practitioners should therefore evaluate prefabrication level together with logistics, factory capacity, site constraints, and robustness evidence.

Recommendation 5: Use robustness evidence before selecting carbon-priority solutions. The Monte Carlo results showed that carbon had the largest uncertainty spread, while the deterministic carbon-minimum representative carried cost and process-hour penalties. Practitioners should therefore avoid selecting a carbon-priority configuration solely from deterministic carbon values unless product-specific carbon data and robustness evidence support the choice.

Recommendation 6: Use the dashboard as a short-listing tool, not as a final design authority.
The dashboard should be used to compare alternatives, inspect S2-relative performance, apply stakeholder weights, and check uncertainty evidence. Final adoption still requires structural verification, project-specific pricing, product-specific carbon assessment, fabricator review, and construction planning.

Table 8-1 Practice recommendations linked to study evidence.

Recommendation	Evidence from thesis	Practical action
Use S2-relative classification	Chapter 5 separated all-objective improvement and trade-off solutions	Classify archive before short-listing
Target moderate rationalisation	Figure 5-8 showed favourable moderate (R_p) and (R_t) zone	Avoid assuming maximum rationalisation is best
Manage opening/detailing complexity	Figure 6-1 and OC-fixed diagnostic showed strong (OC) influence	Review openings, lintels, jambs, trimming, and fit-up burden early
Check prefabrication trade-offs	Figures 4-6, 5-9, and Chapter 6 robustness results showed factory–site effects	Evaluate (P_f) with logistics and site-hour consequences
Use robustness before carbon-priority selection	Table 6-5 showed carbon uncertainty and carbon-minimum trade-off behaviour	Check Monte Carlo probabilities before selecting carbon-led solutions
Use dashboard for short-listing	Chapter 7 workflow defined dashboard-supported interpretation	Use dashboard to compare, rank, and short-list; do not treat it as final approval

8.6: Future Research

Future work should focus on five targeted improvements. First, the framework should be tested on additional CFS residential projects with different geometry, opening layouts, roof forms, and prefabrication strategies. Second, the coefficient set should be recalibrated using measured fabrication time, site installation records, contractor pricing, and product-specific carbon data where available. Third, the OC variable should be developed into a more direct quantity-derived index using opening count, opening area, lintel and jamb details, non-standard panel count, cut-length diversity, and local connection/detailing indicators. Fourth, coefficient-level Monte Carlo samples should be retained so that formal contribution-to-variance or tornado analysis can identify the dominant sources of cost, carbon, and process-hour uncertainty. Fifth, the dashboard should be tested with CFS fabricators, designers, or contractors to evaluate whether the workflow supports actual early-stage decision-making.

8.7: Closing Statement

The thesis demonstrates that early-stage prefabricated CFS framing decisions can be screened more transparently when cost, A1–A5 carbon, process-hours, rationalisation, detailing complexity, prefabrication level, uncertainty, and stakeholder priorities are considered together. The main contribution is not a new optimisation algorithm or a final design tool, but a coefficient-audited and robustness-aware decision-support workflow for identifying and interpreting promising CFS framing configurations in the New Zealand residential context. The results support the use of the framework as a practical short-listing method before detailed engineering, commercial, carbon, and construction review.

Chapter 9: References

1. AbouHamad, M., & Abu-Hamd, M. (2019). Framework for construction system selection based on life cycle cost and sustainability assessment. *Journal of Cleaner Production*, *241*, 118397. https://doi.org/10.1016/j.jclepro.2019.118397

2. Aghasizadeh, S., Tabadkani, A., Hajirasouli, A., & Banihashemi, S. (2022). Environmental and economic performance of prefabricated construction: A review. *Environmental Impact Assessment Review*, *97*, 106897. https://doi.org/10.1016/j.eiar.2022.106897

3. American Iron and Steel Institute. (2022). *North American Specification for the Design of Cold-Formed Steel Structural Members* (ANSI/AISI S100-16(R2020) w/S3-22; 2022 Edition). American Iron and Steel Institute.

4. Anil Kumar, M. V., & Kalyanaraman, V. (2018). Interaction of Local, Distortional, and Global Buckling in CFS Lipped Channel Compression Members. *Journal of Structural Engineering*, *144*(2), 04017192. https://doi.org/10.1061/(ASCE)ST.1943-541X.0001935

5. Audet, C., Bigeon, J., Cartier, D., Le Digabel, S., & Salomon, L. (2021). Performance indicators in multiobjective optimization. *European Journal of Operational Research*, *292*(2), 397–422. https://doi.org/10.1016/j.ejor.2020.11.016

6. Ballesteros-Pérez, P., Kabiri, S., Smith, S. T., & Hughes, W. (2017). Dealing with weather-related claims in construction contracts: A new approach. *Proceedings of International Structural Engineering and Construction*, *4*(1). https://doi.org/10.14455/ISEC.res.2017.221

7. Bataglin, F. S., Viana, D. D., Formoso, C. T., & Bulhões, I. R. (2020). Model for planning and controlling the delivery and assembly of engineer-to-order prefabricated building systems: Exploring synergies between Lean and BIM. *Canadian Journal of Civil Engineering*, *47*(2), 165–177. https://doi.org/10.1139/cjce-2018-0462

8. Blismas, N., Pasquire, C., & Gibb, A. (2006). Benefit evaluation for off-site production in construction. *Construction Management and Economics*, *24*(2), 121–130. https://doi.org/10.1080/01446190500184444

9. BRANZ. (2025). *LCAQuick: Life cycle assessment tool*. https://www.branz.co.nz/environment-zero-carbon-research/framework/lcaquick/

10. Brown, G., Sharma, R., & Kiroff, L. (2020). *Insights into the New Zealand prefabrication industry*.

11. Chaturvedi, S., Rajasekar, E., & Natarajan, S. (2020). Multi-objective building design optimization under operational uncertainties using the NSGA II algorithm. *Buildings*, *10*(5), 88. https://doi.org/10.3390/buildings10050088

12. Coello, C. A. C., Pulido, G. T., & Lechuga, M. S. (2004). Handling multiple objectives with particle swarm optimization. *IEEE Transactions on Evolutionary Computation*, *8*(3), 256–279. https://doi.org/10.1109/TEVC.2004.826067

13. Coello Coello, C. A, Lamont, G. B, & Van Veldhuizen, D. A. (2007). *Evolutionary algorithms for solving multi-objective problems*. Springer US. https://doi.org/10.1007/978-0-387-36797-2

14. *Cost comparison studies*. (2026). Steelconstruction.Info. https://steelconstruction.info/Cost_comparison_studies

15. Deb, K. (2001). *Multi-objective optimization using evolutionary algorithms*. J. Wiley.

16. Deb, K. (2004). *Multi-objective optimization using evolutionary algorithms* (Repr). Wiley.

17. Deb, K., & Jain, H. (2014). An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part I: Solving problems with box constraints. *IEEE Transactions on Evolutionary Computation*, *18*(4), 577–601. https://doi.org/10.1109/TEVC.2013.2281535

18. Deb, K., Pratap, A., Agarwal, S., & Meyarivan, T. (2002). A fast and elitist Multi-objective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, *6*(2), 182–197. https://doi.org/10.1109/4235.996017

19. Doctolero, L., & Batikha, M. (2018). *Using cold-formed steel section in buildings-comparative study*.

20. Doukari, O., Seck, B., & Greenwood, D. (2022). The creation of construction schedules in 4D BIM: A comparison of conventional and automated approaches. *Buildings*, *12*(8), 1145. https://doi.org/10.3390/buildings12081145

21. European Committee for Standardization. (2011). *Sustainability of construction works—Assessment of environmental performance of buildings—Calculation method* (EN 15978:2011). CEN. https://www.cencenelec.eu/news-events/news/2026/en-in-the-spotlight/2026-04-17-en-15978-2026/

22. European Committee for Standardization (CEN). (2026). *Sustainability of construction works—Assessment of environmental performance of buildings—Requirements and guidance* (EN 15978:2026). CEN. https://standards.iteh.ai/catalog/standards/cen/6200b514-96a9-4d5f-9e8f-ca74592f6aea/en-15978-2026

23. Evins, R. (2013). A review of computational optimisation methods applied to sustainable building design. *Renewable and Sustainable Energy Reviews*, *22*, 230–245. https://doi.org/10.1016/j.rser.2013.02.004

24. Ferdous, W., Bai, Y., Ngo, T. D., Manalo, A., & Mendis, P. (2019). New advancements, challenges and opportunities of multi-storey modular buildings – A

state-of-the-art review. *Engineering Structures*, *183*, 883–893. https://doi.org/10.1016/j.engstruct.2019.01.061

25. Few, S. (2006). *Information dashboard design: The effective visual communication of data* (1st ed). O'Reilly.
26. Forsythe, P. J., & Sepasgozar, S. M. E. (2019). Measuring installation productivity in prefabricated timber construction. *Engineering, Construction and Architectural Management*, *26*(4), 578–598. https://doi.org/10.1108/ECAM-09-2017-0205
27. *FRAMECAD*. (2026). https://framecad.com
28. Gadelhak, M., Lang, W., & Petzold, F. (2017). *A visualization dashboard and decision support tool for building integrated performance optimization*.
29. Grubb, P. J. (2001). *Building design using cold formed steel sections: Light steel framing in residential construction*. Steel Construction Institute.
30. Hamdan, M., & Qudah, O. (2015). The initialization of evolutionary multi-objective optimization algorithms. In Y. Tan, Y. Shi, F. Buarque, A. Gelbukh, S. Das, & A. Engelbrecht (Eds.), *Advances in Swarm and Computational Intelligence* (Vol. 9140, pp. 495–504). Springer International Publishing. https://doi.org/10.1007/978-3-319-20466-6_52
31. Hancock, G. J., Chen, Z., Xie, Y., & Pham, C. H. (2026). Recent developments in the direct strength method localised loading design of cold-formed steel sections. In *Recent Advances in Thin-Walled Structures* (pp. 37–51). Elsevier. https://doi.org/10.1016/B978-0-443-36438-9.00004-8
32. Heer, J., & Shneiderman, B. (2012). Interactive dynamics for visual analysis. *Communications of the ACM*, *55*(4), 45–54. https://doi.org/10.1145/2133806.2133821
33. Helton, J. C., & Davis, F. J. (2003). *Latin Hypercube Sampling and the Propagation of Uncertainty in Analyses of Complex Systems*.
34. Hopfe, C. J., & Hensen, J. L. M. (2011). Uncertainty analysis in building performance simulation for design support. *Energy and Buildings*, *43*(10), 2798–2805. https://doi.org/10.1016/j.enbuild.2011.06.034
35. Hwang, C.-L., & Yoon, K. (1981). *Multiple attribute decision making* (Vol. 186). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-48318-9
36. Ian Page & David Norman. (2024). *Prefabrication and standardisation potential in buildings* (STUDY REPORT SR 312 (2014)). BRANZ.
37. Iman, R. L., & Conover, W. J. (1982). A distribution-free approach to inducing rank correlation among input variables. *Communications in Statistics - Simulation and Computation*, *11*(3), 311–334. https://doi.org/10.1080/03610918208812265

38. Inventory of Carbon and Energ. (2026). *ICE*. Embodied Carbon - The ICE Database. Circular Ecology. https://circularecology.com/embodied-carbon-footprint-database.html

39. IPCC. (2023). Buildings. In *Climate Change 2022—Mitigation of Climate Change* (1st ed., pp. 953–1048). Cambridge University Press. https://doi.org/10.1017/9781009157926.011

40. Ishibuchi, H., Masuda, H., Tanigaki, Y., & Nojima, Y. (2015). Modified distance calculation in generational distance and inverted generational distance. In A. Gaspar-Cunha, C. Henggeler Antunes, & C. C. Coello (Eds.), *Evolutionary Multi-Criterion Optimization* (Vol. 9019, pp. 110–125). Springer International Publishing. https://doi.org/10.1007/978-3-319-15892-1_8

41. ISO 14040:2006. (2006). Iso: 14040: 2006. *International Organization for Standardization (ISO), Switzerland*. https://scholar.google.com/scholar_lookup?title=Environmental%20Management%E2%80%93Life%20Cycle%20Assessment%E2%80%93Principles%20and%20Framework&author=ISO%2C%20I.&publication_year=2006

42. ISO 14044:2006. (2006). *ISO 14044:2006*. https://www.iso.org/obp/ui/en/#iso:std:iso:14044:ed-1:v1:en

43. Jaillon, L., Poon, C. S., & Chiang, Y. H. (2009). Quantifying the waste reduction potential of using prefabrication in building construction in Hong Kong. *Waste Management*, *29*(1), 309–320. https://doi.org/10.1016/j.wasman.2008.02.015

44. Lam, K. C., Lee, D., & Hu, T. (2001). Understanding the effect of the learning–forgetting phenomenon to duration of projects construction. *International Journal of Project Management*, *19*(7), 411–420. https://doi.org/10.1016/S0263-7863(00)00025-9

45. Li, Z., Shen, G. Q., & Xue, X. (2014). Critical review of the research on the management of prefabricated construction. *Habitat International*, *43*, 240–249. https://doi.org/10.1016/j.habitatint.2014.04.001

46. Liang, H., Roy, K., Fang, Z., & Lim, J. B. P. (2022). A critical review on optimization of cold-formed steel members for better structural and thermal performances. *Buildings*, *12*(1), 34. https://doi.org/10.3390/buildings12010034

47. *Life Cycle Stages | One Click LCA*. (2025, March 19). https://help.oneclicklca.com/en/articles/275901-life-cycle-stages

48. Lin, C. D., & Tang, B. (2022). *Latin hypercubes and space-filling designs*. https://doi.org/10.48550/ARXIV.2203.06334

49. Liu, H., Al-Hussein, M., & Lu, M. (2015). BIM-based integrated approach for detailed construction scheduling under resource constraints. *Automation in Construction*, *53*, 29–43. https://doi.org/10.1016/j.autcon.2015.03.008

50. Lu, W., Lee, W. M. W., Xue, F., & Xu, J. (2021). Revisiting the effects of prefabrication on construction waste minimization: A quantitative study using bigger data. *Resources, Conservation and Recycling*, *170*, 105579. https://doi.org/10.1016/j.resconrec.2021.105579
51. Lu, W., Yang, Z., & Kong, L. (2023). Identification of learning effects in modular construction manufacturing. *Automation in Construction*, *154*, 105010. https://doi.org/10.1016/j.autcon.2023.105010
52. Lui Brame. (2024). *Waikato Housing Initiative 2023 Housing Stocktake Update*.
53. Mao, C., Xie, F., Hou, L., Wu, P., Wang, J., & Wang, X. (2016). Cost analysis for sustainable off-site construction based on a multiple-case study in China. *Habitat International*, *57*, 215–222. https://doi.org/10.1016/j.habitatint.2016.08.002
54. McKay, M. D., Beckman, R. J., & Conover, W. J. (1979). A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, *21*(2), 239. https://doi.org/10.2307/1268522
55. Mckay, M. D., Beckman, R. J., & Conover, W. J. (2000). A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, *42*(1), 55–61. https://doi.org/10.1080/00401706.2000.10485979
56. Ministry of Business, Innovation and Employment. (2020a). *Building for Climate Change: Transforming the Building and Construction Sector to reduce emissions and improve climate resilience*.
57. Ministry of Business, Innovation and Employment. (2020b). *Whole-of-life embodied carbon framework* (Building for Climate Change Programme) [Consultation Framework / Government Report]. Ministry of Business, Innovation and Employment.
58. Montgomery & Douglas C. (2017). *Design and Analysis of Experiments* (9th ed.). Wiley.
59. Morris, M. D., & Mitchell, T. J. (1995). Exploratory designs for computational experiments. *Journal of Statistical Planning and Inference*, *43*(3), 381–402. https://doi.org/10.1016/0378-3758(94)00035-T
60. Munzner, T. (2014). *Visualization analysis and design* (0 ed.). A K Peters/CRC Press. https://doi.org/10.1201/b17511
61. Navid, A., Khalilarya, S., & Abbasi, M. (2018). Diesel engine optimization with multi-objective performance characteristics by non-evolutionary nelder-mead algorithm: Sobol sequence and latin hypercube sampling methods comparison in DoE process. *Fuel*, *228*, 349–367. https://doi.org/10.1016/j.fuel.2018.04.142
62. *New Zealand Green Building Council*. (2024). http://www.nzgbc.org.nz/

63. Nguyen, A.-T., Reiter, S., & Rigo, P. (2014). A review on simulation-based optimization methods applied to building performance analysis. *Applied Energy*, *113*, 1043–1058. https://doi.org/10.1016/j.apenergy.2013.08.061

64. N.Usefi, Sharafi, P., Mortazavi, M., Ronagh, H., & Samali, B. (2021). Structural performance and sustainability assessment of hybrid-cold formed modular steel frame. *Journal of Building Engineering*, *34*, 101895. https://doi.org/10.1016/j.jobe.2020.101895

65. *NZ Construction Costs Update—QV CostBuilder*. (2025). https://costbuilder.qv.co.nz/nz-construction-costs-update/

66. Onat, N. C., & Kucukvar, M. (2020). Carbon footprint of construction industry: A global review and supply chain analysis. *Renewable and Sustainable Energy Reviews*, *124*, 109783. https://doi.org/10.1016/j.rser.2020.109783

67. Pan, W., Gibb, A. G. F., & Dainty, A. R. J. (2008). Leading UK housebuilders' utilization of offsite construction methods. *Building Research & Information*, *36*(1), 56–67. https://doi.org/10.1080/09613210701204013

68. Pomponi, F., & Moncaster, A. (2018). Scrutinising embodied carbon in buildings: The next performance gap made manifest. *Renewable and Sustainable Energy Reviews*, *81*, 2431–2442. https://doi.org/10.1016/j.rser.2017.06.049

69. *Procurement guide to reducing carbon emissions in building and construction: A practical guide*. (2021). New Zealand Government Procurement.

70. Qingfu Zhang & Hui Li. (2007). MOEA/D: A multiobjective evolutionary algorithm based on decomposition. *IEEE Transactions on Evolutionary Computation*, *11*(6), 712–731. https://doi.org/10.1109/TEVC.2007.892759

71. Saltelli, A., & Annoni, P. (2010). How to avoid a perfunctory sensitivity analysis. *Environmental Modelling & Software*, *25*(12), 1508–1517. https://doi.org/10.1016/j.envsoft.2010.04.012

72. Schafer, B. W. (2008). Review: The direct strength method of cold-formed steel member design. *Journal of Constructional Steel Research*, *64*(7–8), 766–778. https://doi.org/10.1016/j.jcsr.2008.01.022

73. Schuldt, S. J., Nicholson, M. R., Adams, Y. A., & Delorit, J. D. (2021). Weather-related construction delays in a changing climate: A systematic state-of-the-art review. *Sustainability*, *13*(5), 2861. https://doi.org/10.3390/su13052861

74. Shahzad, W. M. (2011). *Offsite manufacturing as a means of improving productivity in New Zealand construction industry: Key barriers to adoption and improvement measures*. Massey University.

75. Shahzad, W. M., Reddy, S. M., Kahandawa, R., & Rotimi, J. O. B. (2024). Benefits, constraints and enablers of modular offsite construction (MOSC) in New Zealand

high-rise buildings. *Engineering, Construction and Architectural Management*, *31*(10), 4042–4061. https://doi.org/10.1108/ECAM-10-2022-1020

76. Shneiderman, B. (1996). *The eyes have it: A task by data type taxonomy for information visualizations*.

77. Song, C., & Kawai, R. (2023). Monte carlo and variance reduction methods for structural reliability analysis: A comprehensive review. *Probabilistic Engineering Mechanics*, *73*, 103479. https://doi.org/10.1016/j.probengmech.2023.103479

78. Song, H., Chen, X., & Wang, H. (2024). Carbon emission optimization for assembled buildings using interval grey GERT modelling and modify NSGA-III algorithm in china. *KSCE Journal of Civil Engineering*, *28*(12), 5415–5426. https://doi.org/10.1007/s12205-024-2619-6

79. Standards Australia / Standards New Zealand. (2018). *AS/NZS 4600:2018* (AS/NZS 4600:2018). https://codehub.building.govt.nz/resources/46002018asnzs#!#resource-detail

80. Stats NZ. (2026). *Building consents issued: December 2025* [ISSN 1178-0231]. Stats NZ. https://www.stats.govt.nz/information-releases/building-consents-issued-december-2025

81. *Steel Frames NZ*. (2026). Https://Steelframesolutions.Co.Nz/. https://steelframesolutions.co.nz/

82. *SteelHaus*. (2024, June 24). https://www.steelhaus.co.nz/get-steel-frame-price/

83. Veljkovic, M., & Johansson, B. (2006). Light steel framing for residential buildings. *Thin-Walled Structures*, *44*(12), 1272–1279. https://doi.org/10.1016/j.tws.2007.01.006

84. Vickers, J., Tan, C. W., Chandrakumar, C., & Vickers, J. (2024). *Emission factors for New Zealand*.

85. Wang, Q., Xu, X., Ding, X., Chen, T., Deng, R., Li, J., & Jiang, J. (2024). Multi objective optimization and evaluation approach of prefabricated component combination solutions using NSGA-II and simulated annealing optimized projection pursuit method. *Scientific Reports*, *14*(1), 16688. https://doi.org/10.1038/s41598-024-65319-3

86. *Whole-of-life embodied carbon assessment: Technical methodology*. (2022). Ministry of Business, Innovation and Employment.

87. Widanage, C., & Kim, K. P. (2024). Integrating design for manufacture and assembly (DfMA) with BIM for infrastructure. *Automation in Construction*, *167*, 105705. https://doi.org/10.1016/j.autcon.2024.105705

88. World Steel Association. (2022). *Life cycle inventory (LCI) study: 2021 data release – Eighth global LCI study for steel products* [Technical Report]. World Steel Association.

89. Xiao, L., Li, Q.-Y., Li, H., & Ren, Q. (2022). Loading capacity prediction and optimization of cold-formed steel built-up section columns based on machine learning methods. *Thin-Walled Structures*, *180*, 109826. https://doi.org/10.1016/j.tws.2022.109826

90. Xu, X., Wang, Q., Ding, X., Chen, T., & Deng, R. (2025). Multi-objective optimization design approach for prefabricated buildings to minimize cost, duration and carbon emissions using ant colony algorithm. *Scientific Reports*, *15*(1), 31683. https://doi.org/10.1038/s41598-025-17216-6

91. Yu, W., & LaBoube, R. A. (2010). *Cold-formed steel design* (1st ed.). Wiley. https://doi.org/10.1002/9780470949825

92. Yu, W., LaBoube, R. A., & Chen, H. (2020). *Cold-formed steel design* (Fifth edition). Wiley.

93. Yuan, Z., Man, Q., Guan, Z., Yi, C., Zheng, M., Chang, Y., & Li, H. X. (2024). Simulation and optimization of prefabricated building construction considering multiple objectives and uncertain factors. *Journal of Building Engineering*, *86*, 108830. https://doi.org/10.1016/j.jobe.2024.108830

94. Zitzler, E., Laumanns, M., & Thiele, L. (2001). *Improving the strength pareto evolutionary algorithm*.

95. Zolghadr, A., Gharaie, E., & Naderpajouh, N. (2022). Barriers to innovation in the housing sector: Economic justifiability of offsite construction for housebuilders. *Journal of Building Engineering*, *52*, 104490. https://doi.org/10.1016/j.jobe.2022.104490

Chapter 10: Appendices

Appendix A: Final Code and Output Documentation

This appendix records the final computational implementation and output archive used to generate the results reported in Chapters 4–7. It supports reproducibility by identifying the final optimisation script, dashboard file, run configuration, output archive, version-control record, and computational environment. It does not introduce new results.

The final optimisation model implementation used for the study was:

cfs_thesis_model.py

This file is the final Python implementation used to generate reported optimisation, sensitivity, robustness, ranking, figure, and dashboard-export outputs.

This model implementation contains the final CFS framing optimisation workflow, including the S2 reference case, decision-variable encoding, objective-function evaluation, LHS-initialised NSGA-II optimisation, Pareto-front extraction, representative solution selection, sensitivity analysis, Monte Carlo robustness analysis, TOPSIS ranking, figure export, and dashboard data export.

The final dashboard file was:

index.html

The dashboard reads the exported optimisation and post-processing outputs and presents them as an interactive decision-support interface. It supports Pareto-front exploration, S2 comparison, representative-solution inspection, sensitivity and robustness review, stakeholder ranking, bounded scenario checks, and model/data audit functions.

The optimiser-controlled design vector used in the final implementation was:

$$\mathbf{x} = [R_p, R_t, OC, P_f]$$

where R_p is the panel rationalisation intensity, R_t is the truss rationalisation intensity, OC is the opening/detailing complexity index, and P_f is the prefabrication level. The derived panel_scheme and truss_scheme labels are reporting categories only and are not independent optimiser-controlled variables.

Table 10-1 Final optimisation and analysis run configuration.

Setting	Final value
Optimisation algorithm	NSGA-II
Initialisation method	Latin Hypercube Sampling
Main seed	42
Multi-seed set	11, 22, 33, 44, 55
Population size	200
Generations	300
Hypervolume samples	500,000
Hypervolume reporting frequency	Every 10 generations
Grid-reference step	0.05
Monte Carlo analysis	Enabled

Setting	Final value
Monte Carlo simulations	1,000
Multi-seed stability check	Enabled
LHS-based sensitivity screening	Enabled
LHS sensitivity samples	2,000

The run manifest was exported as run_manifest_{ts}.csv and records the run timestamp, run mode, seeds, population size, generations, Monte Carlo settings, sensitivity settings, and output folders.

The final run generated three main output locations. These folders were archived together with the final script and dashboard file.

Table 10-2 Final output archive structure.

Folder / file	Purpose
cfs_Thesis_model.py	Final Python optimisation and analysis script
index.html	Final interactive dashboard file
results_fixed/tables	Timestamped CSV outputs for optimisation, verification, diagnostic, sensitivity, uncertainty, ranking, and audit results
results_fixed/figures	PNG figures used in Chapters 5–7
dashboard/data	Dashboard-ready JSON and CSV files
run_manifest_{ts}.csv	Record of run settings, timestamp, and output folders
environment_record_{ts}.txt	Python and package-version record
README_reproducibility.txt	Optional archive guide explaining file structure and how to reproduce outputs

The detailed list of all CSV and PNG files is retained in the archived results folder and repository README rather than repeated in full in the study text.

GitHub Repository and Version-Control Record:

The final Python script and dashboard implementation were archived in a GitHub repository to support reproducibility and traceability. The repository provides a version-controlled record of the final implementation files used for the study .

The GitHub repository is used as a reproducibility archive for the final script, dashboard implementation, and supporting output structure. The numerical results reported in Chapters 5 to 7 remain tied to the final timestamped run and archived output folder. Later repository edits should be treated as subsequent versions and should not be used to reinterpret the submitted thesis results unless a new timestamped run is generated, archived, and documented.

Repository: https://kevinramani-sudo.github.io/cfs-framing-decision-support/dashboard/

The submitted study results should therefore be traced first to the archived run outputs, with the repository used as supporting reproducibility documentation.

All reported study outputs traceable to one final timestamped run. The timestamp links the Pareto-front table, representative solutions, Monte Carlo robustness outputs, component breakdowns, figures, and dashboard-ready files.

Final study run timestamp: 20260527_142125

The computational environment was recorded in environment_record_{ts}.txt and include the Python version, operating system, and exact package versions used in the final run.

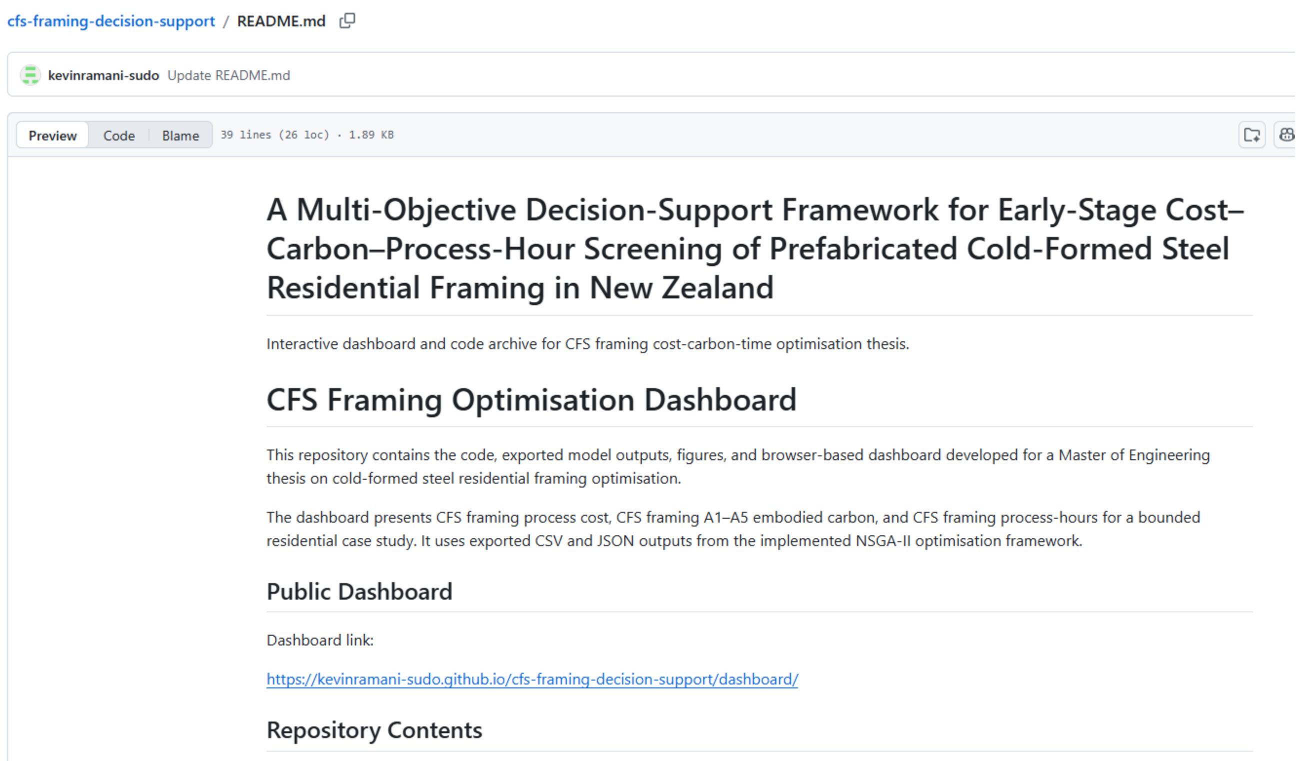

cfs-framing-decision-support / README.md

kevinramani-sudo Update README.md

Preview Code Blame 39 lines (26 loc) · 1.89 KB

A Multi-Objective Decision-Support Framework for Early-Stage Cost–Carbon–Process-Hour Screening of Prefabricated Cold-Formed Steel Residential Framing in New Zealand

Interactive dashboard and code archive for CFS framing cost-carbon-time optimisation thesis.

CFS Framing Optimisation Dashboard

This repository contains the code, exported model outputs, figures, and browser-based dashboard developed for a Master of Engineering thesis on cold-formed steel residential framing optimisation.

The dashboard presents CFS framing process cost, CFS framing A1–A5 embodied carbon, and CFS framing process-hours for a bounded residential case study. It uses exported CSV and JSON outputs from the implemented NSGA-II optimisation framework.

Public Dashboard

Dashboard link:

https://kevinramani-sudo.github.io/cfs-framing-decision-support/dashboard/

Repository Contents

Figure 10-1 GitHub repository used to archive the final optimisation script and dashboard implementation.

Appendix B: Dashboard Implementation Guide

This appendix provides reproducible instructions for opening and using the interactive dashboard developed for this study . The dashboard is a decision-support interface that presents the exported optimisation, sensitivity, uncertainty, ranking, validation, and S2 reference outputs in an interactive format. It does not introduce new results; it provides an accessible interface for exploring the final model outputs reported in Chapters 5–7.

The dashboard archive should contain the following files and folders:

Table 10-3 Required dashboard files.

File or folder	Purpose
index.html	Main interactive dashboard interface
dashboard/data/dashboard_config.json	Main dashboard configuration file
dashboard/data/latest_manifest.json	Record of latest exported dashboard data files
dashboard/data/*.csv	Exported optimisation, sensitivity, uncertainty, ranking, diagnostic, verification, and S2-comparison tables
README.md	Short instruction file for opening the dashboard
results_fixed/tables	Archived timestamped CSV outputs
results_fixed/figures	Archived study figures generated from the final model run

The dashboard distributed with the same timestamped result set used in the study. If new model outputs are generated later, they are treated as separate run and should not be mixed with the reported study results.

The dashboard can be opened using the archived HTML file and exported data folder provided with the study submission. A hosted access copy may also be provided through the project GitHub Pages link. The archived thesis output folder remains the authoritative source for the reported results; the hosted dashboard is used only to support access, reproducibility, and interpretation of the exported model outputs.

The dashboard can be opened using the archived HTML file and exported data folder.

The recommended procedure is:

1. Download or clone the final GitHub repository or study archive.
2. Confirm that index.html and the dashboard/data folder are in the expected relative locations.
3. Open index.html in a modern web browser.
4. If automatic data loading is restricted by the browser, run a local server from the dashboard folder using:

python -m http. server 8000

Then open:

http://localhost:8000/index.html

5. Confirm that the dashboard loads the final exported data, including the Pareto front, representative solutions, Monte Carlo robustness table, sensitivity results, TOPSIS ranking, and S2 reference values.
6. Use the dashboard tabs to explore the results and export a decision or scenario memo where required.

The dashboard is organised into pages that correspond to the main interpretation tasks required in the study.

Table 10-4 Dashboard page structure.

Dashboard page	Primary data	Main decision-support function
Design Explorer	S2 reference values, response model, component data	Inspect selected design-variable settings and compare cost, carbon, and process-hours against S2
Pareto Explorer	Pareto front, representative solutions, parallel-coordinate data	Filter and interpret nondominated solutions by cost, carbon, time, (R_p), (R_t), (OC), (P_f), and derived scheme categories
Scenario Assumptions / Recalibration Check	Dashboard scenario controls, baseline assumptions	Stress-test selected cost, carbon, labour, transport, and emission-factor assumptions within bounded ranges
Sensitivity & Robustness	OAT sensitivity, Spearman/PRCC sensitivity, Monte Carlo robustness	Identify influential variables and assess uncertainty robustness relative to S2
Decision Ranking	TOPSIS stakeholder ranking, robustness penalty settings	Compare preferred alternatives under cost-priority, carbon-priority, time-priority, and balanced decision weights
Model & Data Audit	S2 values, verification/diagnostic status, imported output status, model-scope notes	Confirm data loading, model boundary, validation status, and interpretation limitations

Dashboard page	Primary data	Main decision-support function
Data Import / Export	CSV/JSON import and export functions	Import updated result tables and export selected decision or scenario summaries

The dashboard should load the exported CSV and JSON files from dashboard/data. Where possible, it should display the number of loaded rows for the Pareto front, representative solutions, sensitivity outputs, Monte Carlo robustness table, ranking table, and validation outputs.

If a file is missing, the dashboard should display a clear warning and continue loading the available data. Missing files should not cause the whole interface to fail. This behaviour is important because the dashboard is a decision-support interface that may be used with different result bundles.

The dashboard includes export functions for decision-support reporting. A selected solution summary or scenario memo should record:

- selected solution label.
- design-variable values.
- cost, carbon, and process-hour objective values.
- comparison against S2.
- stakeholder ranking information, where available.
- Monte Carlo robustness information, where available.
- scenario assumptions used in the dashboard.
- model boundary and interpretation limitations.

The exported summary is intended for design discussion and study documentation. It is not a certified design report, cost estimate, or compliance document.

The dashboard and exported summaries are intended for study documentation and early-stage design discussion. They do not replace project-specific structural verification, supplier pricing, certified carbon assessment, fabrication-capacity review, site-constructability assessment, or professional judgement.

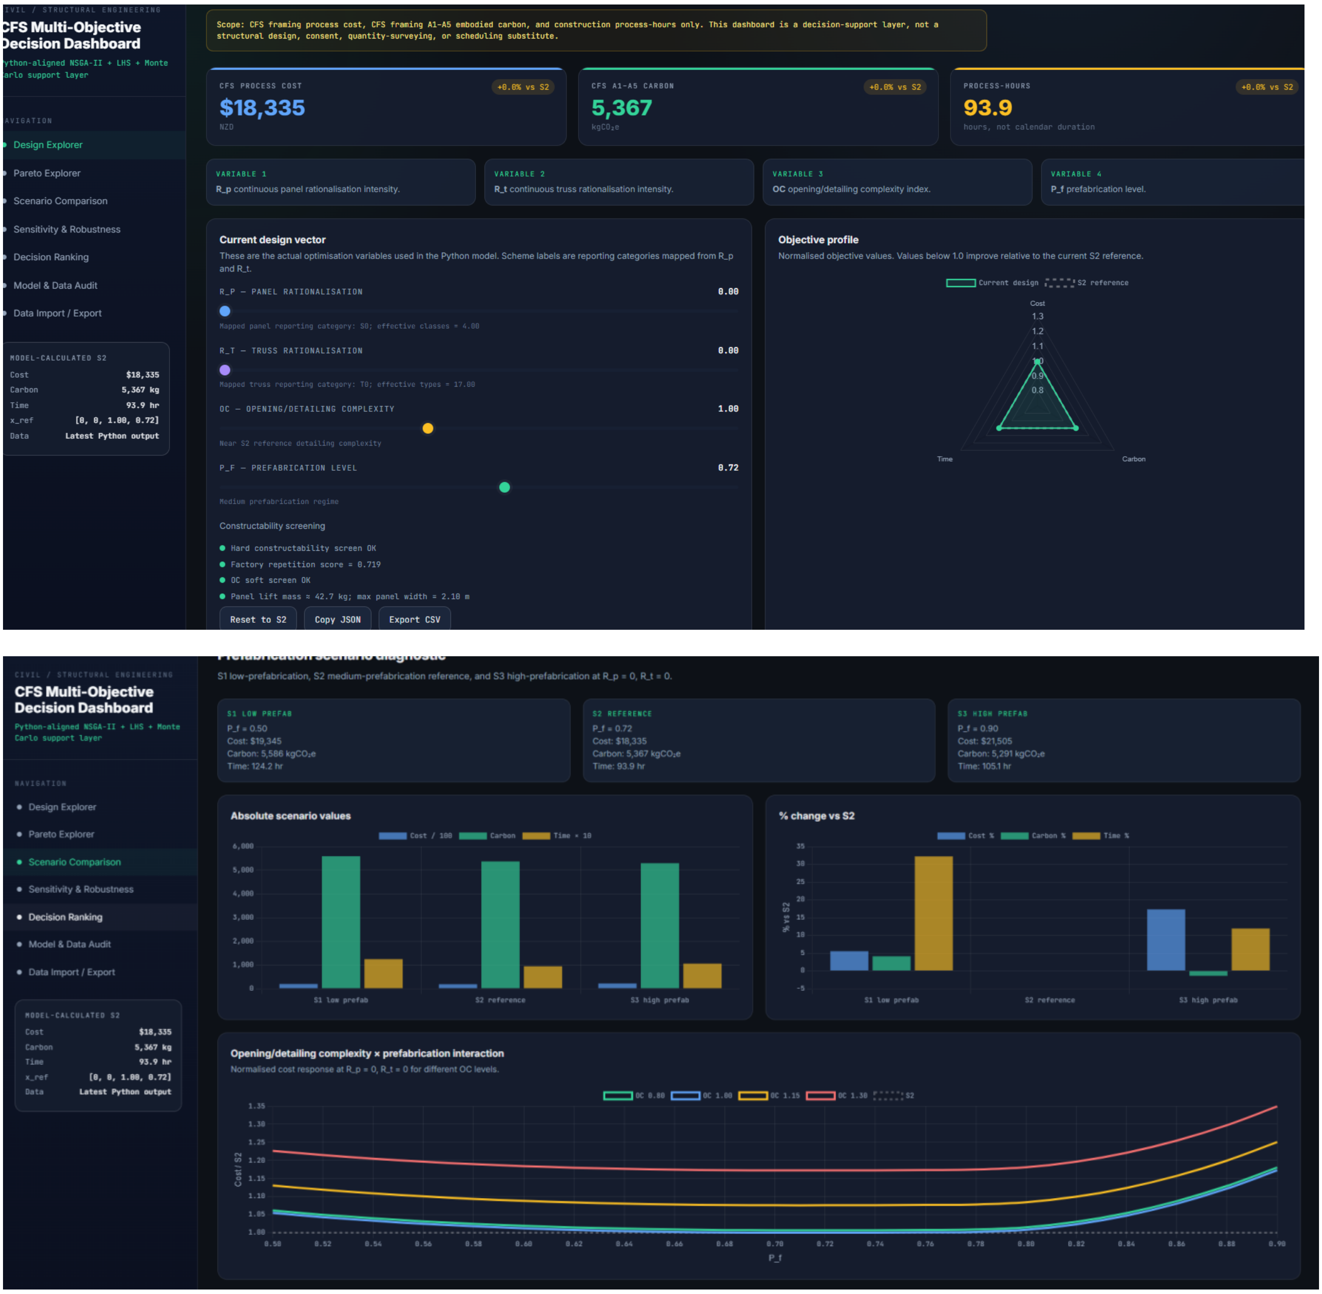

Figure 10-2 Dashboard Pages and Decision-Support Functions.