

Bounded Rationality in Construction-Phase Change Impact Assessment

Mechanisms shaping scope
and depth

Y.A. Oosterlee

MSc Construction Management and Engineering

Delft University of Technology



Bounded Rationality in Construction-Phase Change Impact Assessment

Mechanisms shaping scope and depth

by

Y.A. Oosterlee

To obtain the degree of Master of Science at the TU Delft,
To be defended publicly on Friday September 11, 2026 at 3:30 PM.

Student number: 4954394
Project duration: March 1, 2026 – September 11, 2026
Thesis committee: Prof. dr. P.W. Chan, TU Delft, Chair
Dr. Ir. S. Van Nederveen, TU Delft, Supervisor
Ir. H.J.M. van Triest, Aryo Systems B.V.

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

Dankwoord

Dit afstudeeronderzoek is de afsluiting van mijn masteropleiding Construction Management & Engineering aan de Faculteit Civiele Techniek en Geowetenschappen van de TU Delft. Het schrijven van een masterthesis over besluitvorming, impactanalyses (CIA) en beperkte rationaliteit in complexe bouwprojecten was een uitdagend en leerzaam proces. Ik had dit eindresultaat niet kunnen behalen zonder de steun, inzichten en betrokkenheid van verschillende mensen, die ik hier graag wil bedanken. Ik wil mijn afstudeercommissie eerst bedanken voor hun waardevolle begeleiding, scherpe blik en nuttige feedback tijdens mijn afstudeertraject:

Prof. dr. P.W. Chan (Chair, TU Delft): Bedankt voor uw heldere theoretische begeleiding, overzicht en het zorgen voor de academische kwaliteit van dit onderzoek. Uw feedback heeft geholpen om de complexe ideeën van begrensde rationaliteit duidelijk en logisch te presenteren.

Dr. Ir. S. Van Nederveen (Supervisor, TU Delft): Dank u voor uw voortdurende steun, uw bereikbaarheid en de waardevolle inhoudelijke feedback. Onze gesprekken gaven steeds de goede richting om de balans te houden tussen de praktijk en de diepere academische kennis.

Ir. H.J.M. van Triest (Aryo Systems B.V.): Heel erg bedankt voor de fijne samenwerking en het bieden van de praktische context bij Aryo Systems B.V. Jouw inzichten in de praktische situatie van veranderingen in de bouwfase, tijdsdruk en dynamiek tussen verschillende partijen waren essentieel om het theoretische model te testen aan echte projecten.

Ik wil ook alle respondenten en professionals bedanken die de tijd hebben genomen om mee te doen aan de interviews en casusreconstructies. Het delen van uw ervaringen uit de praktijk, openhartige reflecties en projectdocumentatie heeft de praktische basis van deze thesis gevormd. Dit onderzoek was niet mogelijk geweest zonder uw bereidheid om openlijk te praten over de problemen die komen kijken bij beslissing nemen onder tijdsdruk.

Ik wil graag mijn familie, vriendin, vrienden en studiegenoten bedanken voor hun constante steun, interesse en de gezellige momenten buiten het studeren. Jullie steun en vertrouwen in de afgelopen maanden hebben mij de kracht gegeven om dit traject succesvol af te maken.

Youp Oosterlee
Delft, september 2026

Abstract

Construction-phase Change Impact Assessment (CIA) is routinely formally complete but remains selectively narrow. Impacts not surfaced before commitment often resurface later as corrective work. This thesis asks how bounded rationality affects the quality of construction-phase CIA and explains how that selective narrowing arises in complex engineering projects.

The framework is empirically tested on six CIA episodes in a single industrial construction complex and triangulated with six semi-structured interviews: two with practitioners involved in the case, two with cross-project procedural experts, and two with practitioners working on adjacent projects. The findings are reported throughout at calibrated evidence weights (Strong, Moderate, Preliminary).

Time pressure is confirmed as the precipitating condition that determines whether an assessment enters a bounded rational regime at all. Information uncertainty and cognitive load determine which mechanism becomes salient within a bounded rational regime. This is the baseline condition, not a variable to be reduced. Four refinements follow. A meta-BR layer at the contracting stage results in an inherited assessment environment through a design-maturity pathway (Strong) and a procurement-fragmentation pathway (Moderate). Early closure operates on two tiers, per change and at portfolio level via bundled settlement, expanding the scope from individual to decision portfolios. Four protective conditions bounded the mechanisms. Three were derived from the framework and one, relational backing, was inductively derived as the meta-condition that supports the others under sustained time pressure. Finally, the three constraints work as a self-reinforcing system rather than as parallel inputs, with each decision reducing the resources available for the next.

Read through an ecological account of heuristics, the mechanisms turn out to be ambivalent. The same strategies that narrowed assessment and corrective work in four episodes also closed two others with none. Bounded rationality thus both supports and undermines assessment quality, depending on whether a protective condition makes the residue of a simplified assessment visible and recoverable.

Three of the six tools are used within a single assessment cycle, and three are used across cycles. A workflow specification carrying the tools into the documentation environment of the case organization is also presented. The results shift CIA's practical goal from completeness to managed incompleteness: knowing what was not assessed and having a plan for what to watch for during implementation.

List of Abbreviations

Core thesis concepts

BR	Bounded rationality
CIA	Change Impact Assessment
EP	Episode; the six analysed episodes are labelled EP-01 to EP-06
INT	Interview; the six semi-structured interviews are labelled INT-01 to INT-06
SQ	Sub-question

Framework codes

A1	Time pressure (triggering constraint)
A2	Information incompleteness (triggering constraint)
A3	Uncertainty (triggering constraint)
B1	Truncated information search (mechanism)
B2	Heuristic substitution (mechanism)
B3	Early decision closure (mechanism), refined into B3a and B3b
B3a	Per-change early closure
B3b	Portfolio-level early closure via bundled settlement
D1–D6	Downstream corrective-work codes (D1 redesign, D2 resequencing, through to D6 no corrective work)
CL	Coding Log reference; CL-024 denotes Coding Log entry 24

Contextual and technical

AHP	Analytic Hierarchy Process
BIM	Building Information Modelling
CBS	Cause Breakdown Structure
CII	Construction Industry Institute
DSM	Design Structure Matrix
FEED	Front-End Engineering Design
FID	Final Investment Decision (project-stage authorisation gate)
FIDIC	Fédération Internationale Des Ingénieurs-Conseils (international contract standards)
IFC	Industry Foundation Classes (BIM data schema)
NAC	National Academy of Construction
RII	Relative Importance Index

Contents

Dankwoord	ii
Abstract	iii
List of Abbreviations	iv
1. Introduction	1
1.1. Context and problem statement	1
1.2. Research objective	2
1.3. Research questions	2
1.4. Research scope	2
1.5. Relevance of the study	3
1.5.1. Scientific relevance	3
1.5.2. Practical relevance	3
1.6. Thesis structure	4
2. Literature Review	5
2.1. Complex engineering projects and construction-phase change	5
2.2. Types and sources of construction-phase changes	6
2.2.1. Stakeholder-driven sources of change	6
2.2.2. Site conditions and external triggers	7
2.2.3. Typical impact categories and why they matter for CIA	8
2.3. Change Impact Assessment (CIA): approaches and limitations	8
2.3.1. Dependency- and trace-based CIA approaches	8
2.3.2. BIM- and graph-based CIA: modeling project topology	9
2.3.3. Uncertainty-aware CIA: taxonomies and fuzzy-logic tool support	9
2.3.4. Multi-criteria and stakeholder-weighted assessment approaches	9
2.3.5. Limitations: why CIA tools underperform during construction	10
2.4. Decision-making under time pressure and uncertainty: bounded rationality	10
2.4.1. Satisficing and the limits of search	10
2.4.2. Two readings of heuristics: bias and ecological fit	11
2.4.3. Three mechanisms of constrained assessment	11
2.4.4. Bounded rationality beyond the single decision	12
2.5. Corrective work as a downstream manifestation	12
2.6. Synthesis and research gap	13
3. Theoretical Framework	15
3.1. Purpose and boundary of the framework	15
3.2. Core assumptions of bounded rationality in construction-phase CIA	15
3.3. Mechanisms: from constraints to CIA limitations	16
3.3.1. Mechanism 1: truncated information search under time pressure	16
3.3.2. Mechanism 2: heuristic substitution through role-based relevance filtering	16
3.3.3. Mechanism 3: early decision closure under unresolved uncertainty	17
3.4. Protective conditions: when mechanisms do not narrow assessment	18
3.5. CIA completeness as an outcome variable: scope and depth	18
3.6. Conceptual model	20
3.7. Analytical implications for methodology	21
4. Methodology	22
4.1. Research approach	22

4.2. Research design	23
4.3. Case selection	23
4.3.1. Applying the selection criteria	23
4.3.2. Contractual context	24
4.3.3. A complex case	24
4.4. Data collection	25
4.4.1. Semi-structured interviews	25
4.4.2. Document analysis	26
4.4.3. Reconstruction of CIA episodes	26
4.5. Data analysis	26
4.5.1. Step 1: within-case analysis	27
4.5.2. Step 2: process tracing	27
4.5.3. Step 3: cross-case comparison	27
4.5.4. Quote selection	27
4.5.5. Mechanism evidence weighting	27
4.6. Quality criteria and validity	28
4.7. Ethical considerations	28
5. Findings	29
5.1. Overview of CIA episodes	29
5.2. CIA as conducted in practice	32
5.2.1. Common process structure	32
5.2.2. Variations that matter analytically	32
5.3. Constraints observed	33
5.3.1. Time pressure	33
5.3.2. Information incompleteness	35
5.3.3. Uncertainty	35
5.3.4. Upstream conditions: contracting-stage origins	36
5.4. Bounded rationality strategies observed	37
5.4.1. Truncated information search (B1)	37
5.4.2. Heuristic substitution (B2)	38
5.4.3. Early decision closure (B3)	39
5.4.4. Protective conditions	39
5.4.5. Institutional patterns beyond the per-change mechanisms	41
5.5. CIA completeness outcomes: scope and depth	43
5.5.1. Scope: breadth of domains, actors, and interfaces considered	44
5.5.2. Depth: reach of impact reasoning	44
5.6. Downstream manifestations: corrective work	44
5.6.1. Episodes with corrective work	44
5.6.2. Contrastive episode: no corrective work	45
5.6.3. Cross-episode pattern	46
6. Discussion	47
6.1. Interpreting findings through bounded rationality	47
6.1.1. Time pressure as the dominant trigger	47
6.1.2. Bounded rationality is partly inherited from the contracting stage	50
6.1.3. Per-change mechanisms: B1 truncated search and B2 heuristic substitution	50
6.1.4. Per-change early closure (B3a) and portfolio-level early closure (B3b)	51
6.1.5. Protective conditions: a fourth layer	51
6.2. Implications for CIA theory	52
6.3. Implications for practice	53
6.3.1. Upstream: front-end design maturity (against meta-BR)	54
6.3.2. Per-change: pre-formal alignment and conditional approval (against B3a)	54
6.3.3. Portfolio: periodic review (around B3b)	54
6.3.4. Cross-cutting: relational investment (sustaining the protective layer)	54
6.4. Limitations	55

7. Practical Framework and Recommendations	56
7.1. A bounded-rational CIA support perspective	56
7.2. Tools inside a single change cycle	58
7.2.1. Tool 1: Depth Check (While the assessment is being written)	58
7.2.2. Tool 2: Question check at the assessment moment (When the consultation closes)	58
7.2.3. Tool 3: Open-items approval (At the approval itself)	59
7.2.4. The three tools applied to EP-01	59
7.3. Standing conditions around every cycle	62
7.3.1. Tool 4: Inherited-gap report (Once per change, read at Director level)	62
7.3.2. Tool 5: Bundle review (Once a quarter)	62
7.3.3. Tool 6: Continuity rule (Once per project)	62
7.3.4. Cross-cycle boundary	63
7.4. Carrying the tools into an operational environment	63
8. Conclusion	65
8.1. Answer to the main research question	65
8.2. Answers to sub-questions	66
8.2.1. SQ1: how is CIA enacted in practice?	66
8.2.2. SQ2: what constraints characterize CIA episodes?	66
8.2.3. SQ3: which strategies are used, and under what conditions?	66
8.2.4. SQ4: how do strategies shape CIA completeness, and how do missed impacts become visible later?	67
8.2.5. SQ5: under what conditions do strategies not reduce CIA completeness?	67
8.3. Theoretical contribution	67
8.4. Practical contribution	68
8.5. Recommendations for future research	69
Bibliography	70
A. Interview Guide	72
Preamble en geïnformeerde toestemming	72
Fase 1 — Algemene procesbeschrijving	72
Fase 2 — Concrete episode	73
Afsluiting	75
Post-interview analytisch memo	75
B. Interview Transcripts	77
C. Coding Log	78
C.1 Evidence-weight summary by framework component	78
C.2 Cross-interview map	79
C.3 New findings register	81
D. Coding Scheme	82
D.1 Constraint codes	82
D.2 Mechanism codes	82
D.3 Protective conditions	83
D.4 Downstream corrective work codes	84
D.5 Context and new candidate codes	84
D.6 Evidence-weight categories	85
D.7 Quote selection criteria	85
E. Overview of Interviews Conducted	86
F. Declaration on the use of artificial intelligence	89

List of Figures

2.1. Taxonomy of construction-phase change sources	7
3.1. CIA completeness matrix: scope and depth	19
3.2. Conceptual model: bounded rationality and CIA completeness	20
4.1. Research design overview	22
4.2. Structure of a reconstructed CIA episode: stages, evidence sources, and analytical focus	26
5.1. CIA episodes positioned on the scope and depth matrix	43
6.1. Refined conceptual model with feedback dynamics	49
7.1. When each intervention fires within the CIA cycle	57
7.2. The three point-of-use tools tested against the six episodes	61

List of Tables

2.1. Overview of CIA approach families, core assumptions, and key limitations in construction-phase contexts	8
5.1. Overview of CIA episodes	31
5.2. Constraint intensity by episode	34
5.3. Bounded rationality strategies by episode	42
5.4. Downstream corrective work by episode	45
C.1. Evidence-weight summary by framework component	78
C.2. Cross-interview map	80
C.3. New findings register	81
D.1. Constraint codes (A)	82
D.2. Bounded rationality mechanism codes (B)	82
D.3. Protective conditions (PC)	83
D.4. Downstream corrective work codes (D)	84
D.5. Context (CTX) and new candidate (NEW) codes	84
D.6. Evidence-weight categories	85
D.7. Quote selection criteria	85
E.1. Overview of interviews conducted	87

1

Introduction

1.1. Context and problem statement

Large engineering projects are often executed in environments that face strong technical interdependence and evolving information. During the construction phase, design intentions are confronted with site-specific realities, while multiple disciplines must coordinate decisions while work is still ongoing. In this context, requirement- and design-related changes are routine and must be evaluated, and decisions must be made within limited time frames. Changes could be the result of shifting stakeholder demands, incomplete or immature design information, unforeseen site conditions, or external contractual and regulatory pressures (Bano et al., 2012; Jayatilleke & Lai, 2018).

The challenge is therefore not the occurrence of change itself, but the ability to anticipate its consequences at the moment decisions are taken. Once construction is underway, interfaces are already established, and work may have progressed beyond the point at which changes can be implemented without disruption. Further corrective action may be needed downstream if pertinent implications are not found prior to implementation. Redesign, rescheduling, demolition, or reconstruction are some examples of this additional corrective action. Rework is a frequent term for this corrective action, which is associated with higher expenses, delays in schedules, and decreased productivity (Forcada et al., 2017; Hwang et al., 2009; Love et al., 2022). However, rework is viewed as a consequence of the shortcomings in the previous change evaluation rather than as the main issue.

Change Impact Assessment (CIA) is used in engineering projects to support choices about proposed changes. This is the procedure for determining and evaluating the consequences of a suggested change prior to its execution (Arnold & Bohner, 1993). The technical dependencies and simulation of the spread of changes between components, interfaces, and specifications are two aspects of many current CIA techniques (Gräßler et al., 2022; Sun et al., 2006). These methods can provide a helpful framework for analysis, but they frequently depend on requirements that are challenging to fulfill during construction. For instance, reliable and easily available system data and enough time for methodical examination. This makes it impossible to formally conclude assessments made during the construction phase while keeping their scope and level of information constrained.

Within the limitations of the construction phase, this thesis conceptualizes the completion of CIA as the result of bounded rationality-based decision-making. The concept of bounded rationality holds that decision-makers make decisions under time constraints, with limited cognitive ability and imperfect knowledge (Kahneman, 2003; Simon, 1955). This leads assessors to rely on practical simplifications: truncated information search, heuristic reasoning, and early decision closure. These simplifications constrain CIA along two dimensions. Scope refers to the disciplines, interfaces, stakeholders, and follow-up activities taken into account. Depth refers to the extent to which the impact analysis extends beyond direct local effects. When scope or depth are limited, certain impacts may remain latent and only become visible later as corrective work.

Concretely, the problem this thesis examines is what professionals do when they must make a change decision on incomplete information. Time pressure is what closes the window of opportunity, but the substantive challenge is informational in nature: the knowledge required to assess a change is scattered across disciplines, organizations, and documents. This information is also rarely consolidated in one place at the time the decision needs to be made. The trend is illustrated by one instance later in this thesis (EP-01, chapter 5). A watercourse crossing had to be reconstructed, a dam had to be replaced

1. Introduction

by a culvert, and a choice had to be made in 23 days to meet a predetermined external date. The examination was efficient and procedurally sound, and the crossing was adequately analyzed. But the body that controlled that watercourse had its own terms, and those terms were just outside the boundaries of the assessment. They were only apparent after they were approved and the design work had to be re-done. The assessment was not negligent, and the missing information was not unknowable; it was outside the boundary the assessors had time to draw. This thesis asks what professionals do at that boundary, what information they pursue, what proxy they substituted, and what they knowingly proceed without.

1.2. Research objective

The aim of this research is to explain how bounded rational decision-making under constraints during the construction phase shapes the scope and depth of Change Impact Assessment. This thesis will not consider CIA as a purely technical evaluation. But treats, CIA as a decision-making episode in which actors must make an evaluation in the face of time constraints, incomplete information, and continuous production demands (Kahneman, 2003; Simon, 1955). It seeks to identify the decision-making processes that lead to assessments becoming selectively limited and to explain how these restrictions enable some effects to go unnoticed during the decision-making process.

In support of this, CIA completeness is graded as a continuum defined by scope (domains and stakeholders included) and depth (impact reasoning over time and interdependencies). This operationalization allows for systematic comparison between CIA episodes and supports mechanism-based explanation rather than descriptive reporting. The study also intends to translate its findings into practical insights: leverage points for improving CIA within realistic construction constraints. It does this without assuming that full evaluation is ever possible while construction is still going on (Love et al., 2022).

1.3. Research questions

Main research question

| *How does bounded rationality influence the quality of construction-phase Change Impact Assessment?*

Sub-questions

SQ1. How is CIA enacted in practice during construction-phase change decisions?

SQ2. What time and information constraints characterize CIA episodes during construction?

SQ3. Which bounded rationality strategies are used during CIA, and under what conditions do they occur?

SQ4. How do these strategies shape CIA completeness (scope and depth), and how do missed impacts become visible later as corrective work?

SQ5. Under what conditions do bounded rationality strategies *not* affect CIA completeness, and what sustains those conditions?

1.4. Research scope

This research focuses on construction-phase CIA episodes in complex engineering projects, defined as the period in which a proposed change is assessed, discussed, and decided while construction activities are ongoing. The study adopts a qualitative multiple-case study design using change cases from projects executed by Aryo Systems in construction-intensive environments. The unit of analysis is the CIA episode rather than the overall project, supporting detailed examination of decision-making under pressure.

The scope is intentionally bounded to changes that occur when ongoing work constrains assessment time and information quality. Specifically, the study considers requirement-driven and design-related changes arising during construction, modifications to specifications, design details, interfaces, or execution requirements (Bano et al., 2012; Jayatilleke & Lai, 2018). Changes originating purely in early design phases are excluded unless they re-emerge and require reassessment during construction.

The unit of analysis remains the construction-phase CIA episode. Where the empirical material indicates that the constraints operating within an episode were established earlier, however, the analysis follows them to their origin. Conditions set at the contracting stage, such as the maturity of the design package at the start of construction or the number of interfaces created by the procurement structure, are therefore treated as part of the explanation of construction-phase CIA rather than as a separate object of study. This is an analytical extension rather than a widening of the case boundary: no pre-construction decision episode is itself reconstructed, and the upstream conditions are evidenced through respondent accounts rather than through documentary analysis of the contracting stage.

Downstream corrective work is examined as an outcome indicator rather than as the primary research object. The study does not aim to quantify overall cost overruns, claims, or lifecycle impacts. Instead, it examines how limitations in CIA scope and depth can allow impacts to remain latent and later require corrective effort, such as redesign, resequencing, demolition, and reconstruction, or additional coordination (Love et al., 2022). The analysis is mechanism-focused and deliberately avoids claims of deterministic causality. Given its qualitative design, the study seeks analytical rather than statistical generalization, aiming to identify recurring decision mechanisms that explain how CIA becomes selectively narrow under construction-phase constraints.

1.5. Relevance of the study

1.5.1. Scientific relevance

Changes are frequent and frequently linked to unfavorable outcomes including rework, productivity loss, and cost overruns, according to research on engineering and construction projects (Forcada et al., 2017; Hwang et al., 2009; Love et al., 2022). Simultaneously, a significant amount of scholarship has created techniques and resources to help CIA predict the effects of suggested modifications (Gräßler et al., 2022; Sun et al., 2006). However, there has not been much integration at the level of fundamental explanatory mechanisms between these two study lines, which have mostly grown independently.

The majority of current CIA research views assessment as a technical and analytical endeavor with an emphasis on developing formal procedures and tools, tracing change propagation, or detecting dependencies (Gräßler et al., 2022; Sun et al., 2006). Although this study explains how impacts could be evaluated in theory, it does not take into account how CIA is actually conducted during construction, where time constraints, uncertainty, and changing information limit what can be evaluated (Kahneman, 2003). Because of this, the research offers little explanation for why, even in cases when formal assessment procedures are in place, CIA frequently maintains limited scope or depth.

By creating a bounded rationality explanation of CIA limits in the construction environment, this thesis takes that gap into account. It views CIA completeness as the result of making decisions under limits and connects these restrictions to corrective work downstream as a likely manifestation of latent repercussions (Love et al., 2022). By doing this, it provides a mechanism-based explanation that enhances knowledge of how change decisions are actually made in practice and supports tool-oriented CIA research.

1.5.2. Practical relevance

Decisions about changes are often made in construction projects under extreme time constraints and with insufficient information. Under such circumstances, CIA tends to concentrate on immediate and visible effects, whereas indirect or downstream repercussions only show up after implementation, when corrective work is more expensive and disruptive (Love et al., 2022). Current change management guidelines often focus on documentation and procedural compliance (Sun et al., 2006), but they offer little assistance for enhancing assessment quality within practical construction limits.

1. Introduction

This study offers practical leverage points for practice by identifying the bounded rationality strategies that systematically narrow CIA scope and depth. The study advocates for context-sensitive enhancements, such as focused reminders, light cross-checks, and explicit stopping criteria, rather than assuming that comprehensive evaluation is possible. Each is made to effectively increase scope and depth in the presence of warning signals, lowering the possibility of needless corrective action downstream while recognizing the limitations practitioners truly encounter.

1.6. Thesis structure

This thesis is structured to develop a mechanism-based understanding of CIA limitations under construction-phase decision constraints. [chapter 1](#) introduces the research context, problem framing, objective, questions, scope, and relevance. [chapter 2](#) reviews literature on construction-phase change, CIA approaches, and decision-making under time pressure and uncertainty, culminating in the identification of the research gap. [chapter 3](#) presents the bounded rationality theoretical framework and conceptual model, specifying the mechanisms through which constraints shape CIA scope and depth. [chapter 4](#) describes the methodology and explains how CIA episodes are reconstructed and analyzed. [chapter 5](#) reports findings on CIA practice, observed constraints, bounded rationality strategies, CIA completeness outcomes, and downstream corrective work. [chapter 6](#) discusses how the findings refine the bounded rationality explanation and relate to existing literature. [chapter 7](#) translates the insights into practical implications and a bounded-rational perspective on CIA support. [chapter 8](#) answers the research questions and sets out the theoretical and practical contributions, the limitations, and directions for future research.

2

Literature Review

This chapter explores the reasons why, in situations where teams are faced with tight deadlines and uncertainty, construction-phase Change Impact Assessments (CIAs) often fall short in scope and depth. System complexity and interdependence in projects, the ambiguity of construction changes, current CIA models and their underlying assumptions, bounded rationality in fast-paced decision-making, and corrective work as a trailing indicator of uncaptured impacts are the five main themes that organize the literature. In the end, this review identifies a significant gap in the literature: how teams are forced to use a selectively narrow lens during impact assessments due to real-time construction constraints.

2.1. Complex engineering projects and construction-phase change

In addition to their size, complex engineering and construction projects are difficult due to the intricate relationships between organizational actors, technical systems, and outside limitations. The degree of interdependence between project attributes and interfaces and the degree to which this interdependence impacts schedule and cost predictability as well as functional performance is a common definition of complexity in project management literature (Construction Industry Institute, 2015; National Academy of Construction, 2022). Because it presents complexity as an information-processing issue, this definition is especially pertinent to this thesis. Projects may still accomplish their intended functionality as interdependencies grow, but they may become more challenging to forecast and manage in terms of time and expense.

The critical consequence is that complex projects cannot be understood entirely through linear decomposition. Although systems are normally depicted as separable components in drawings, specifications or work packages, dependencies exist across interfaces so that changes to a local element can have indirect and non-obvious effects elsewhere in the system (Clarkson et al., 2004; Jarratt et al., 2011). Complexity is therefore socio-technical, as it stems not only from the technical architecture of the system but also from the distribution of knowledge, responsibilities and decision authority between several organizations and teams (Lafhaj et al., 2024; National Academy of Construction, 2022). As the number of stakeholders and coordination boundaries increases, the relevant information about dependencies becomes more and more fragmented and difficult to assemble within limited decision windows.

The literature also indicates that project complexity is multi-dimensional rather than driven by a single factor. Reviews highlight the recurring complexity drivers such as the number of stakeholders, project scope and objectives, management structure, design and technological novelty, and regulatory uncertainty (Lafhaj et al., 2024). Industry frameworks also emphasize governance arrangements, stakeholder alignment, interface density, resource constraints and execution targets (Construction Industry Institute, 2015; National Academy of Construction, 2022). This is critical to construction-phase change, because each of these drivers influences both the location of impact knowledge, and the ease of retrieval, verification and coordination of that knowledge in the time available to support a decision.

The construction phase is particularly prone to complexity related disruption as it is where design intent meets site conditions under ongoing production pressure. Construction has been described as a complex and dynamic system where the interactions between crews, equipment, material flows and partially completed systems produce results that are difficult to predict from isolated task descriptions

(National Academy of Construction, 2022). Interdependencies are operationally binding during this phase, interface assumptions are translated into physical constraints, sequencing dependencies are constraining flexibility, and late information can require rapid reconfiguration of planned work. Recent work on large-scale infrastructure projects further demonstrates the use of dependency representations for predicting change propagation, but also points to the difficulty of capturing the full system context under real project constraints (Chen & Whyte, 2022). Therefore, changes during construction generally propagate from the immediate change to downstream activities, constructability and cross-disciplinary coordination needs (Clarkson et al., 2004).

For CIA, this construction-phase context substantially increases the information-processing burden in two ways. First, the range of potentially affected elements expands, as interdependencies span multiple disciplines, contractors, documents, and work packages, and as second- and higher-order effects become more plausible in tightly coupled systems (Clarkson et al., 2004; Jarratt et al., 2011). Second, the assessment object itself becomes less stable: proposed changes may enter assessment while still evolving, and information about current site conditions, progress status, and interface readiness is often incomplete or dispersed across actors (Lafhaj et al., 2024). As a result, CIA in the construction phase resembles a time-constrained decision episode rather than a bounded technical calculation, requiring cross-boundary search, interpretation, and coordination under uncertainty (National Academy of Construction, 2022; Woods & Dekker, 2000). This provides the task environment in which bounded rationality mechanisms are likely to arise, and it is this connection that the following sections develop.

2.2. Types and sources of construction-phase changes

A change in the construction phase is common in project delivery because execution reveals differences between design assumptions, contractual scope, and on-site conditions. In construction management literature it is useful to distinguish between the deviation, which is a change in the work, and the administrative instrument that formalizes the deviation. A change order is the formal contractual record to authorize and price a deviation (Keane et al., 2010). A variation is a deviation from the originally agreed scope, quantities, specifications or sequence of work. This difference is relevant to CIA because the subject of assessment is usually the proposed variation, which is often incomplete or continuing to evolve, whereas the change order is the consequence of negotiation and decision closure under time pressure.

Another difference is the manner in which changes are introduced in the project and what this means for impact reasoning. Some studies differentiate between variation orders that may yield value improvement through optimization or quality improvement and those that yield value degradation through errors, omissions or poor planning that require rework or disruption. However, the distinction between beneficial and detrimental changes is often unclear at the time of decision, as downstream effects may not be visible within the assessment window (Al Amri & Marey-Pérez, 2020; Ibbs et al., 1998). The important implication for this thesis is not that changes can be neatly labeled in advance. Instead, differences in change origins affect urgency, information availability, and the probability of indirect effects, all of which affect the scope and depth of CIA under bounded rationality.

Construction-phase changes do not arise from a single source but from the interaction of multiple stakeholder actions, site conditions, and external pressures. Figure 2.1 organizes these sources into two broad categories: stakeholder-driven changes, originating from owners, designers and consultants, or contractors; and site- or externally-driven changes, triggered by unforeseen ground conditions, access constraints, or regulatory and supply-chain disruptions. This distinction matters for CIA because each source carries a different urgency profile, information availability, and dependency structure, all of which condition how far impact reasoning can extend within a given decision window.

2.2.1. Stakeholder-driven sources of change

Empirical studies have repeatedly pointed to the actions and interactions of core stakeholders, owners, consultants and designers, and contractors, who control different information and constraints, as being responsible for changes during construction. Recent work suggests structuring causes by a hierarchical cause breakdown structure (CBS) and emphasizes that change drivers are interrelated rather than




Figure 2.1.: Taxonomy of construction-phase change sources

independent (Muhsen & Rezouki, 2022). This view strengthens a view that is consistent with the complexity literature reviewed in section 2.1: a change is rarely owned by a single party as its consequences propagate across technical interfaces and organizational boundaries.

Owner-related changes often arise from evolving requirements, late scope decisions, financial constraints, or approval delays. These drivers tend to enter CIA with high urgency and contractual pressure, which can compress the time available to explore second- and third-order effects (Keane et al., 2010; Muhsen & Rezouki, 2022). In CIA terms, this increases the likelihood of a truncated search and early closure because the decision is framed around immediate feasibility and cost or time implications rather than broader interface impacts.

Consultant and design-related changes are frequently linked to incomplete, inconsistent, or poorly coordinated design information. Keane et al. (2010) discuss how design development, errors and omissions, and coordination breakdowns can generate variations and disputes, particularly when construction sequencing has already committed resources and interfaces. Muhsen and Rezouki (2022) further show that stakeholders may disagree on what counts as a root cause, for example on bill-of-quantities accuracy or estimation quality, indicating that CIA may start from contested information and divergent interpretations. When problem ownership and data reliability are disputed, assessments tend to privilege accessible cues such as constructability and short-term schedule protection over comprehensive cross-disciplinary propagation reasoning.

Contractor-related changes are often associated with execution constraints such as resource availability, workmanship, means-and-methods adaptation, or procurement disruptions. These changes may be framed as a practical necessity in the field and therefore evaluated under strong production pressure, which again narrows the feasible scope and depth of CIA (Muhsen & Rezouki, 2022). In such cases, the assessment problem is not only technical, what changes, but also organizational: who can supply reliable impact information fast enough. This increases the risk that impacts remain latent until downstream work reveals them as corrective actions.

2.2.2. Site conditions and external triggers

Beyond stakeholder actions, construction-phase change is frequently triggered by site-specific and external conditions that were uncertain or unknowable at earlier stages. Unforeseen ground conditions, spatial constraints, and operational access limitations can force redesign, resequencing, or temporary works that were not included in the original plan (Keane et al., 2010). External triggers, such as regulatory changes or supply-chain disruptions, can similarly compress decision windows and reduce the option space, increasing reliance on rapid judgment rather than systematic impact tracing. For CIA, the key mechanism is that uncertainty here is not merely missing data: it often concerns the stability of the situation itself, as conditions continue to evolve. This makes it difficult to define assessment boundaries and to verify propagation paths before committing to implementation.

2.2.3. Typical impact categories and why they matter for CIA

Recurring domains of impact associated with changes across sources include cost growth, schedule delay, quality risk, and relationship or claims escalation. Standard construction management texts note that changes can increase overhead and disruption costs and create coordination burdens that are higher than the direct line-item value of the changed work. Further, quantitative change research suggests a negative effect of higher amounts of change on productivity and overall project performance (Ibbs et al., 1998). It also indicates possible cumulative effects when changes are interacting over time.

For this thesis, these impact domains are not only outcomes but also search targets in CIA. Assessors under time pressure focus on direct and immediately visible impact, local cost and near-term schedule, and omit indirect impacts that require cross-disciplinary consultation and longer-horizon reasoning (Kahneman, 2003). The type and source of a change, then, pre-structure CIA by determining where relevant knowledge is located, how contested it is, and how costly it is in time and coordination to expand the scope and depth of assessment.

2.3. Change Impact Assessment (CIA): approaches and limitations

The purpose of CIA is to prevent local modifications, drawing revisions or a requirement change from having unanticipated knock-on effects on technical interfaces, tasks and organizational responsibilities. The CIA is needed because project interdependencies are not easily visible. A project topology consists of interdependencies between requirements, building components and construction tasks. This is typically not an explicit design deliverable and therefore opaque to project teams (Isaac & Navon, 2009).

One of the key drivers for CIA is that impacts at the time of non-recognition may surface later as downstream remediation work. However, reported rework and corrective work cost levels vary significantly between studies, partly due to different definitions and measurement boundaries (Forcada et al., 2017; Love et al., 2022). The takeaway for this thesis is not that CIA can eliminate downstream effort, but that CIA is a risk reduction mechanism: increasing the breadth (scope) and reach (depth) of impact reasoning can reduce the chance that important impacts remain latent at the time decisions are closed.

In the literature on CIA, there have emerged a number of methodological approaches, each based on different assumptions about the availability of information, the representation of the system and the analytical capacity of decision-makers. Four main approach families are summarized in Table 2.1: traceability-oriented, BIM and graph-based, uncertainty-aware, and multi-criteria stakeholder methods. The table highlights the main tools associated with each family, the key assumption behind their application and the main limitation they face in the construction phase.

Table 2.1.: Overview of CIA approach families, core assumptions, and key limitations in construction-phase contexts

Approach family	Core assumption	Key limitation
Traceability-oriented	Stable, accessible information	Links missing or outdated in practice
BIM- and graph-based	Complete model of project topology	False confidence
Uncertainty-aware	Fuzzy logic, causal taxonomies, tool-its	Dynamic site context outpaces model updates
Multi-criteria	Stakeholder weights can be elicited	Contestation of weights under time pressure

2.3.1. Dependency- and trace-based CIA approaches

Change effects are transferred through connected elements rather than local (Clarkson et al., 2004; Jarratt et al., 2011) and much CIA research addresses the representation and the tracing of dependencies. The CIA approaches are characterized by the representation of the system and the search

for affected elements in the traditions of systems and software (Arnold & Bohner, 1993). In engineering and construction, two families of approaches are of particular relevance: (a) traceability-oriented approaches that connect requirements, specifications, design elements and deliverables so that an initiating change can be traced forwards and backward across artifacts, and (b) dependency-oriented approaches that model coupling between components, interfaces and tasks to estimate propagation beyond direct links (Arnold & Bohner, 1993; Clarkson et al., 2004).

In construction settings, these approaches are often operationalized using representations such as dependency graphs, traceability matrices, or design structure matrices (DSM), which provide a structured way to identify potentially affected elements and second-order effects (Sosa et al., 2004; Sullivan et al., 2023). Their fundamental contribution is systematization. They reduce the dependence of impact search on memory and informal coordination by encoding interdependencies into an artifact that can be queried.

2.3.2. BIM- and graph-based CIA: modeling project topology

A more construction specific stream tries to mine dependency information from Building Information Models (BIM) and express it as a graph, where interdependencies can be analyzed algorithmically. The topology of a project is defined by Isaac and Navon (2013) as the network of interdependencies of requirements, building components and construction tasks. They propose the synthesis of BIM and graph theory for the analysis of the topology of a project. In their framework, different aspects of the project can be represented as nodes connected through relationships, requirements, and spaces (IFC entities such as `IfcSpace` and `IfcPropertySet`), design components (`IfcElement`) and planned tasks (`IfcTask`) into a single graph model.

The benefit of this approach is that it enables an algorithmic search for the elements that are hit. Shortest-path searches can be used to identify chains of requirements to downstream cost or task elements; breadth-first search can provide an estimate of a local blast radius surrounding a changed component; and clustering can identify tightly coupled subsystems that may be susceptible to propagation (Isaac & Navon, 2009). This is consistent with the wider change-propagation literature in complex design where propagation can follow many paths and the ability to predict likely paths can help to support prioritization of assessment effort (Clarkson et al., 2004).

2.3.3. Uncertainty-aware CIA: taxonomies and fuzzy-logic tool support

Several construction-oriented CIA toolkits integrate causal taxonomies with uncertainty handling. This combination reflects that assessments during the construction phase often start from incomplete or qualitative information. Sun et al. (2006) describe a change management toolkit which includes a change dependency framework for relating project attributes to change causes and effects, and uses fuzzy logic to model imprecise relationships using linguistic scales where precise numerical data are not available. This is important for CIA because many variables that matter for impact, such as site complexity, team experience and quality of coordination, are hard to parameterize but still influence outcomes of change. The key point is that such approaches do not remove uncertainty; they make it explicit and thus discussable. Their relevance when considering this thesis is that they can support good enough CIA by prompting broader scope and deeper reasoning without the need of exhaustive modeling.

2.3.4. Multi-criteria and stakeholder-weighted assessment approaches

The CIA is not merely technical: the perceived seriousness of a change is based on stakeholder priorities, cost, time, quality, safety and operational performance. Some approaches use multi-criteria decision methods to aggregate and prioritize impact dimensions across stakeholder perspectives. For example, Gunduz and Mohammad (2020) use Relative Importance Index (RII) and Analytic Hierarchy Process (AHP) to rank critical success factors, stressing that there may be different opinions between stakeholders on what is most important. The value of such methods for CIA is not that they compute the right decision but that they structure trade-offs and make weighting explicit which is convenient when time pressures force prioritization. From the bounded rationality perspective, explicit criteria and

weights can help mitigate ad hoc heuristic substitution by offering a clear rationale for what is prioritized and what is postponed.

2.3.5. Limitations: why CIA tools underperform during construction

However, despite methodological advances, CIA performance is subject to the availability of information, the completeness of the model, and the ability to process and coordinate information within the time available, conditions that are difficult to meet during construction (Sun et al., 2006; Woods & Dekker, 2000). Three especially pertinent limitations exist.

First, hidden, partial relationships. BIM-focused approaches often contain rich object data, but relational dependency information may be hidden and not utilized unless it is explicitly extracted and represented (Isaac & Navon, 2009). Likewise, traceability and dependency models rely on disciplined data capture. If links are missing or outdated, CIA might appear to be formalized, but it is not complete in practice.

Second, the dynamic site environments and the adaptability gaps. Many CIA representations are for relatively stable design contexts. Construction phase reality is dynamic. Site constraints, access conditions, sequencing commitments and late information can change the effective dependency structure faster than models are updated (Sun et al., 2006; Woods & Dekker, 2000). Tools, thus, can find out what is connected, but cannot yet support timely decision-making if the change object is changing and the decision windows are short.

Third, model simplification and tacit knowledge limitations. Graph-based models intentionally simplify complex interactions into binary/weighted links. This is explicitly mentioned as a strong approximation by Isaac and Navon (2013), as many interactions depend on time, space and method details that are hard to encode. Parametric and rule-based approaches always face the problem that tacit knowledge of experts has to be translated into explicit relationships that can be read by machines. This simplification carries the risk of false confidence: models may underrepresent nuances in the means and methods of construction and therefore under-estimate indirect impacts.

Current CIA approaches taken together offer a useful analytical framework for depicting and investigating interdependencies. However, many implicitly assume stable and accessible information and sufficient time for systematic analysis, assumptions often violated in construction phase change decisions. This is the motivation for the thesis focus on CIA as bounded rational decision-making. Under time pressure and uncertainty, assessments may be formally completed, yet be selective in scope and depth.

2.4. Decision-making under time pressure and uncertainty: bounded rationality

The limitations of existing CIA approaches identified in [section 2.3](#) are not simply a matter of inadequate tools or procedures. They reflect a more fundamental condition: decision-makers conducting CIA during construction cannot comprehensively search all available information, evaluate all possible impacts, or select an optimal course of action. They act instead under bounded rationality. This section develops that lens in four steps: the classical account of satisficing and search ([subsection 2.4.1](#)); two competing readings of what heuristics are and when they work ([subsection 2.4.2](#)); the three mechanisms through which constrained assessment becomes observable ([subsection 2.4.3](#)); and the extension of bounded rationality beyond the individual decision to organizations and contracts ([subsection 2.4.4](#)).

2.4.1. Satisficing and the limits of search

Bounded rationality was introduced by Herbert Simon as a critique of the classical model of the fully rational agent. Simon (1955) argued that real decision-makers do not maximize utility across all available options; instead, they satisfice, searching until they find an option that is good enough relative to some aspiration level and stopping there. Satisficing is not a failure of reasoning but an adaptive response to the irreducible constraints of real decision environments: it allows action to continue without indefinite deliberation.

Two features of Simon's account matter for this thesis. The first is that the binding constraint is informational rather than purely temporal. What limits the decision-maker is not the clock as such, but the cost of acquiring and processing information relative to the value of acquiring more of it; time pressure matters because it raises that cost, not because it is a constraint of a separate kind. The second is that bounded rationality is a claim about the relation between a decision-maker and a task environment, not about the decision-maker alone. Simon (1956) framed rational behavior as shaped jointly by the capabilities of the actor and the structure of the environment in which the actor operates, an argument he later expressed as two blades of a pair of scissors, neither of which cuts alone (Simon, 1990). This relational framing is what makes it possible to ask, as this thesis does, not only whether an assessment was constrained but whether the constrained assessment was well or badly matched to the situation it faced.

2.4.2. Two readings of heuristics: bias and ecological fit

Simon's account has developed into two research traditions that read the same phenomenon in opposite directions, and the distinction between them is consequential for how CIA limitations should be interpreted.

The first is the heuristics-and-biases tradition associated with Kahneman and Tversky. In Kahneman (2003) dual-process formulation, System 1 thinking is fast, automatic and associative, producing rapid judgments from pattern recognition and prior experience, while System 2 is slow, deliberate and analytical, but demands significant effort and time. Under time pressure and uncertainty, System 1 processes dominate. Kahneman and Frederick (2002) specify the central operation as attribute substitution: when confronted with a difficult target question, the decision-maker unconsciously answers an easier heuristic question in its place and treats the answer as though it addressed the original. This tradition explains why assessment under pressure privileges familiar, visible and immediately interpretable impacts over indirect, unfamiliar or longer-horizon consequences, and it measures heuristics against a normative standard of full analysis, treating departures from that standard as bias.

The second tradition, associated with Gigerenzer and colleagues, reads the same behavior as ecologically rational. On this account heuristics are not degraded approximations of optimal reasoning but simple strategies drawn from an adaptive toolbox, and their performance depends on the fit between the structure of the heuristic and the structure of the environment in which it is used (Gigerenzer & Goldstein, 1996; Gigerenzer & Selten, 2001). A heuristic is therefore not good or bad in itself; it is good or bad relative to an environment. Under genuine uncertainty, as distinct from calculable risk, simple strategies that deliberately ignore information can outperform more complex ones, because complex methods overfit noise that simple methods pass over (Gigerenzer & Brighton, 2009; Gigerenzer & Gaissmaier, 2011). Love et al. (2023) apply this framing precisely to construction, exploring how fast-and-frugal heuristics can be assembled into an adaptive toolbox for assessing the uncertainty of rework on a major infrastructure project.

The two traditions are not straightforwardly reconcilable, and this thesis does not attempt to adjudicate between them. They are retained here because they answer different questions, both of which the empirical material requires. The heuristics-and-biases reading explains the mechanism: what is substituted for what, and why the substitution is cognitively attractive under pressure. The ecological reading explains the variance: why the same mechanism, applied in two different episodes, can produce an adequate outcome in one and a costly one in the other. Existing CIA research has implicitly adopted the first reading alone, treating incomplete assessment as a shortfall to be corrected by better tooling. Adopting the second alongside it reframes the practical question from how to eliminate bounded rational strategies, which is not achievable under construction conditions, to how to identify the environments in which each strategy is and is not well matched, and to make that match more reliable.

2.4.3. Three mechanisms of constrained assessment

Drawing these traditions together, three mechanisms can be specified through which constrained decision-making becomes observable in an assessment episode. They are presented here as a synthesis for the purposes of this study rather than as an established taxonomy, and each is anchored in the literature reviewed above.

2. Literature Review

The first is truncated information search. Search stops once the assessment is judged sufficient to support an immediate decision rather than continuing until all plausible impacts have been explored. This follows directly from Simon (1955) stopping rule: search terminates when the marginal cost of continuing exceeds the expected marginal benefit, and under construction-phase conditions that threshold is reached early, because extending search requires coordination across documents, disciplines and organizational boundaries.

The second is heuristic substitution. A difficult assessment question, such as what the full cross-disciplinary impact of a change will be, is replaced by a more accessible proxy question, such as what the change means for the schedule or whether it resembles a previous case, and the answer to the proxy is treated as an answer to the original. This is attribute substitution as specified by Kahneman and Frederick (2002), operating on assessment questions rather than on the judgement tasks studied experimentally.

The third is early decision closure. The assessment is closed while material uncertainty remains unresolved, because the costs of delay, in disrupted production or contractual exposure, are judged to outweigh the benefits of further deliberation. In Simon's terms this is satisficing against a deadline-adjusted aspiration level: the decision-maker accepts a residual of unresolved questions as the price of acting within the available window.

The three mechanisms do not describe individual failure or poor judgement. They are predictable responses to a task environment that demands action faster than comprehensive analysis can be completed, and, following subsection 2.4.2, whether a given instance proves costly is not determined by the mechanism itself but by how well it is matched to the situation and by what, if anything, catches its residue. chapter 3 specifies how the three operate in the construction-phase CIA context and how their effects can be traced empirically.

2.4.4. Bounded rationality beyond the single decision

Simon's account was formulated at the level of an individual decision, but its most consequential extensions concern organizations and contracts, and both are relevant to a setting in which assessment is distributed across roles and governed by contractual instruments.

March and Simon (1958) established that organizations respond to the bounded rationality of their members structurally, through standard operating procedures, divided attention and routinized decision rules. These devices economize on individual cognitive effort, but they also determine in advance which information will be gathered and which questions will be asked, so that the boundaries of an assessment are partly set by organizational routine rather than by the substance of the change under assessment.

Williamson (1975, 1985) developed the contractual extension by pairing bounded rationality with opportunism. Because parties cannot foresee and specify every future contingency, contracts are necessarily incomplete; because counterparties may exploit that incompleteness to their advantage, the incompleteness is not merely a technical gap but a strategic resource. Williamson further showed that where assets are specific to a relationship, the parties become bilaterally dependent and the cost of exiting or renegotiating rises sharply. Construction-phase change decisions sit squarely in this configuration: once work is underway, interfaces are established and progress is physically committed, so the assets are highly specific and lock-in is close to complete at the moment changes must be assessed.

Two implications for this study are following. First, the incompleteness of the contractual framework for change decisions should be anticipated, not seen as a defect in a particular contract, and it is the space that this leaves open that is the site of discretion in assessment practice. Second, in this context bounded rationality need not be self-imposed. If one party knows the decision window of the other, the constraint can be strategically imposed, extending the analysis from the cognitive limits of an assessor to the interaction of contracting parties.

2.5. Corrective work as a downstream manifestation

Corrective work, often labeled rework, is commonly defined as unnecessary effort associated with redoing an activity that was not correct the first time (Hwang et al., 2009; Love et al., 1999). It is associated with cost growth, productivity loss, and schedule disruption (Forcada et al., 2017; Love et al., 2022). In this thesis, corrective work is not treated as the primary problem definition. Instead,

it serves as a downstream indicator that certain impacts were not fully anticipated at the moment of change decision-making.

A useful distinction in recent literature concerns avoidable versus unavoidable corrective work. Some corrective effort could plausibly have been reduced through better anticipation, coordination, or decision timing; other effort reflects genuine uncertainty and emergent interactions that are difficult to predict, regardless of assessment quality (Love et al., 2022). This distinction aligns with a bounded rationality view because it shifts analytical attention from blaming those who made the wrong call to the structural limits of decision-making under constraints. Not all downstream corrective work indicates a failure of CIA, but some of it plausibly does, and identifying which is an empirical question.

The key link to CIA is that corrective work can be interpreted as a delayed manifestation of latent impacts. When CIA scope is narrow, relevant stakeholders or interfaces are excluded, or CIA depth is shallow, impact reasoning is truncated before indirect effects are explored, and some consequences remain hidden until downstream tasks encounter the affected conditions. Empirical work on the social organization of errors reinforces that breakdowns often become visible only after commitments have been made and work has progressed. This thesis, therefore, treats corrective work not as a direct measure of CIA quality, but as a plausible signal that the assessment window closed before certain impacts were surfaced.

2.6. Synthesis and research gap

The literature reviewed in this chapter indicates that construction-phase CIA is challenged by three interacting conditions. First, dense technical interdependencies and complex change propagation expand the information space that CIA must cover; second- and third-order effects become more plausible as coupling increases (Clarkson et al., 2004; Jarratt et al., 2011). Second, changes often emerge during active production with incomplete or still-evolving definitions, making the assessment object itself unstable at the moment assessment begins (Sun & Meng, 2009). Third, decision windows are short, and information is unevenly distributed across actors, disciplines, and documents, making a comprehensive search cognitively and organizationally demanding within the available time.

CIA research has responded primarily through technical and procedural tools: dependency models, traceability links, BIM-graph approaches, fuzzy causal toolkits, and multi-criteria prioritization frameworks (Sosa et al., 2004; Sullivan et al., 2023). These approaches make a valuable contribution by systematizing impact search and making interdependencies explicit. However, as [section 2.3](#) demonstrated, they share a common implicit assumption: that information is sufficiently stable and accessible, and that time is available for systematic analysis. Both assumptions are routinely violated during construction-phase change decisions.

The rework and corrective work literature documents the frequency and consequences of downstream effort and increasingly attributes such outcomes to constrained decision-making, role ambiguity, and communication breakdowns (Love et al., 2022). However, CIA itself is rarely examined as a bounded decision episode. The question of how time pressure, information incompleteness, and uncertainty translate into specific assessment strategies, and how those strategies produce observable, systematic limitations in what is assessed, has received limited empirical attention.

The resulting gap is both theoretical and empirical, and it has three components. The first concerns mechanism: no established account connects construction-phase constraints to specific bounded rational strategies and traces how those strategies constrain CIA completeness along the dimensions of scope and depth. The second concerns variance. If, as [subsection 2.4.2](#) argues, the value of a simplifying strategy depends on its fit to the environment in which it is used, then the more informative question is not whether construction-phase CIA is constrained, which it invariably is, but under what conditions constrained assessment nonetheless proves adequate. The CIA literature offers little on the conditions that contain the consequences of incomplete assessment, as distinct from procedures intended to prevent incompleteness in the first place. The third concerns scale and locus. Bounded rationality in this setting has been examined almost exclusively as a property of an individual assessment episode, leaving open whether the same mechanisms operate across portfolios of changes, whether they are shaped by conditions established before the construction phase begins, and whether they can be induced in one party by another under the contractual arrangements reviewed in [subsection 2.4.4](#).

2. Literature Review

Empirically, there is limited evidence on how these mechanisms operate across different CIA episodes, which configurations of constraints tend to activate which strategies, which conditions distinguish episodes where constrained assessment proved adequate from those where it did not, and which scope and depth limitations are most consistently associated with downstream corrective work.

This thesis addresses these dimensions of the gap. [chapter 3](#) develops a theoretical framework that specifies the mechanisms through which construction-phase constraints translate into selectively narrow CIA, using bounded rationality as the explanatory lens (Kahneman, 2003; Simon, 1955). [chapter 4](#) and [chapter 5](#) then operationalize and empirically examine these mechanisms across CIA episodes from construction projects, enabling mechanism-based explanation rather than descriptive reporting of CIA limitations.

3

Theoretical Framework

3.1. Purpose and boundary of the framework

This chapter develops the theoretical framework explaining why Change Impact Assessment (CIA) can become selectively narrow during the construction phase of complex engineering projects. The framework is grounded in bounded rationality, which conceptualizes decision-making as purposeful but constrained by limited time, incomplete information, and limited cognitive capacity (Kahneman, 2003; Kahneman & Frederick, 2002; Simon, 1955, 1956). The aim is to identify and explain how constraints during the construction phase affect the assessment method, and how this subsequently limits the completeness of the CIA.

The framework is deliberately bounded in scope. It does not attempt to build a full organizational theory of projects, nor does it prescribe an optimal assessment process. Instead, it provides a decision-mechanism explanation aligned with the main research question: how does bounded rationality influence the quality of construction-phase Change Impact Assessment? Following section 1.2, quality is not compared with maximal completeness; the framework specifies both the mechanisms through which assessment narrows and, in section 3.4, the conditions under which a narrowed assessment may nonetheless prove adequate.

Complexity, fragmented information, and multi-actor coordination, reviewed in chapter 2, are treated here as features of the task environment that increase information demands rather than as separate theoretical lenses.

The comprehensiveness of a CIA is a scale with two dimensions: scope and depth. Scope refers to the domains, interfaces, and stakeholders included. Depth concerns the extent of the analysis: does it limit itself to direct, local effects, or does it also consider long-term and chain effects? These two dimensions are developed further in section 3.5, where they are operationalized as an outcome variable for empirical analysis.

3.2. Core assumptions of bounded rationality in construction-phase CIA

Bounded rationality departs from the classical assumptions of comprehensive search and optimal decision-making. Rather than evaluating all available options and selecting the best, decision-makers operating under bounded rationality satisfice. They accept an outcome that is good enough relative to some aspiration level and stop searching once they reach it (Simon, 1955). This is not irrational behavior; it is a practical way to deal with situations where a full analysis is simply impossible because there is too little time and cognitive capacity for it.

In construction-phase CIA, this adaptive condition is the norm rather than the exception. Decision-makers face three recurring constraints. First, time is limited by production schedules and the need to maintain workflow continuity; decisions must be made within windows defined by site activities, not by analytical readiness. Second, information is incomplete and often evolving: uncertainty concerns not only unknown facts but also the reliability and timing of available inputs, and the change object itself may still be shifting when assessment begins. Third, cognitive capacity is finite, meaning attention and

analytical effort must be allocated selectively rather than distributed equally across all possible impacts (Kahneman, 2003).

The three constraints do not operate independently. They form a feedback system in which each constraint amplifies the others, especially under sustained pressure. Time pressure compresses the information-gathering phase, increasing uncertainty at the moment of decision. The combined burden of short time and elevated uncertainty raises cognitive load because the actor must track not only what is known but also what might be missed. The mechanism responses that follow, truncated search, heuristic substitution, and early closure, close decisions with residual uncertainty, and that residue consumes additional time downstream when it later surfaces as corrective work. The constraint set is therefore best understood as a self-reinforcing process rather than as a static triple of independent inputs. This dynamic is developed empirically in [chapter 5](#) and interpretively in [subsection 6.1.1](#).

Under these conditions, the strategies specified in [subsection 2.4.3](#) become predictable responses rather than individual failings. Decision-makers satisfice by accepting an assessment that is actionable within the available window; they truncate their information search before all impacts are explored; and they substitute difficult judgments, 'who might be affected across all disciplines?', with more accessible cues, such as formal role designation. The practical consequence is that CIA tends to become selectively narrow: constrained in scope and depth in ways that allow certain impacts to remain latent at the moment of decision closure. The following section specifies the mechanisms through which this narrowing occurs.

3.3. Mechanisms: from constraints to CIA limitations

The framework specifies three bounded-rationality mechanisms through which construction-phase constraints translate into limitations on CIA completeness. Each mechanism is named, explained as an adaptive response to constraints, and linked to its primary effect on CIA scope and depth.

3.3.1. Mechanism 1: truncated information search under time pressure

During construction-phase CIA, decision-makers frequently operate under tight decision windows imposed by ongoing production activities and sequencing commitments. When available time is limited, information search is truncated once the assessment is judged sufficient to support an immediate decision, rather than continued until all plausible impacts are explored (Simon, 1955). Extending the search for additional impacts, particularly second- and third-order effects, requires time-consuming coordination across documents, disciplines, and organizational boundaries. When such effort threatens to delay execution, decision-makers prioritize completing the assessment over expanding its analytical reach.

Truncated search systematically constrains CIA depth. While first-order and near-term impacts may be identified, indirect consequences related to downstream sequencing, interface behavior, or longer time horizons are less likely to be explored before decision closure. These deeper impacts are not necessarily unknown; they remain undiscovered within the available assessment window. Consequently, they may only become apparent later in implementation, when construction activities encounter conditions not fully anticipated, leading to corrective work such as resequencing, redesign, or additional coordination.

[chapter 5](#) traces this mechanism concretely in EP-01 and EP-04, where formal consultation was broad, up to nine disciplinary roles in EP-04, but the substantive content of what was returned was thin: several inputs consisted of little more than 'no comments'. The professionals involved were not skipping the search; they were asking a smaller number of people a smaller question than full coverage would require, because that was what the decision window allowed.

3.3.2. Mechanism 2: heuristic substitution through role-based relevance filtering

During construction-phase conditions, decision-makers often face real uncertainty about the full range of possible impacts linked to a proposed change. When time pressure and uncertainty limit the feasibility of comprehensive system-wide analysis, decision-makers substitute a difficult question, 'who might

be affected downstream, and how?', with a more accessible cue: formal role designation (Kahneman & Frederick, 2002). Rather than identifying relevance based on the content of actors' knowledge, assessors rely on organizational roles as a proxy for determining whose input is required.

This role-based heuristic is an adaptive response to bounded rationality. Formal roles provide an easily accessible and socially accepted foundation for making consultation decisions, especially when extending the consultation process requires significant time and coordination. However, in complex engineering projects, impact-relevant knowledge is often distributed across actors whose formal responsibilities do not include change assessment: site engineers, supervisors, or subcontractors with situated execution insight. When these actors fall outside the formally defined CIA process, their perspectives are less likely to be mobilized, even if they hold critical knowledge about interfaces, sequencing constraints, or constructability implications.

As a result, heuristic substitution via role-based filtering systematically constrains the CIA's scope. Assessments prioritize impacts aligned with formally represented disciplines and decision domains, while effects that depend on cross-boundary or informal expertise are more likely to be excluded. These excluded impacts are not necessarily less significant; they are less visible within the time-pressured assessment context. Therefore, certain interface-related or execution-dependent impacts remain latent at decision closure and may surface later as corrective work or additional coordination.

chapter 5 and chapter 6 trace several concrete forms this takes: substituting a prior, superficially similar case for a fresh assessment (EP-04, EP-05), substituting the question 'does this fit the milestone envelope?' for 'what is the full impact?' (EP-04), and, in the interview material, substituting relational judgment ('pick your battles') or presenter confidence for verification. In each case, the professional is not failing to reason; they are reasoning with a cheaper, more accessible question because the full question cannot be answered in the time available.

3.3.3. Mechanism 3: early decision closure under unresolved uncertainty

Construction-phase CIA often begins from a change definition that is incomplete, evolving, or contested across stakeholders. Uncertainty may concern the technical solution, the boundaries of the change, site conditions, or the reliability of available input information. When such uncertainty cannot be resolved within the decision window imposed by production demands, decision-makers may proceed to close the assessment despite unresolved questions.

Early decision closure is a bounded rational response to uncertainty under time pressure. Rather than delaying action until greater clarity is achieved, which may be infeasible during active construction, decision-makers accept residual uncertainty and commit to a course of action. This allows work to continue but limits the opportunity to revisit assumptions or explore alternative impact pathways once commitments are made.

Early closure constrains both CIA scope and depth simultaneously. In terms of scope, unresolved uncertainties may discourage the inclusion of additional stakeholders or domains whose input could further complicate the decision. In terms of depth, impact reasoning is curtailed when uncertainty becomes difficult to manage within the available time. Certain impacts, therefore, remain latent at decision closure and become visible only later, when evolving conditions or downstream tasks expose mismatches between assumptions and reality. These latent impacts frequently manifest as corrective work during implementation.

EP-05 shows this mechanism in its most consequential form: a decision closed under a hard production deadline before a municipal ruling was known, protecting the schedule but exposing the project to the full downstream cost once the ruling went the other way. EP-03 shows the same mechanism producing no downstream cost at all, because the closure was made explicitly conditional on the missing information arriving later. The difference between these two outcomes, developed in chapter 6, is not whether the professionals involved acted under uncertainty; they all did, but whether the closure named what remained unresolved.

3.4. Protective conditions: when mechanisms do not narrow assessment

The three mechanisms specified in [section 3.3](#) describe how construction-phase constraints translate into narrowed assessment. They do not, on their own, predict that every constrained assessment will prove inadequate. Following the ecological reading developed in [subsection 2.4.2](#), a simplifying strategy is neither sound nor unsound in itself: its adequacy depends on the fit between the strategy and the environment in which it is applied (Gigerenzer & Brighton, 2009; Gigerenzer & Gaissmaier, 2011). If that is correct, then the framework requires a second class of elements alongside the mechanisms, conditions that improve the fit, or that catch the residue when the fit is poor. These are termed protective conditions.

Three can be derived directly from the structure of the mechanisms themselves, one for each. Each addresses the specific way in which its corresponding mechanism can go wrong, rather than attempting to suppress the mechanism, which the constraints of the construction phase make unrealistic.

The first is institutional slack. Truncated search ([subsection 3.3.1](#)) fails when search stops before impact-relevant knowledge has been reached, and the stopping point is set by the perceived cost of continuing relative to the time available. Any condition that lowers that cost or extends the window therefore raises the threshold at which search terminates. Institutional slack refers to resource or schedule headroom held within the decision process itself, a phased rather than compressed approval architecture, or an interval deliberately preserved between provisional and definitive commitment, which allows search to continue past the point at which production pressure alone would have ended it.

The second is the preparatory heuristic. Heuristic substitution ([subsection 3.3.2](#)) fails not because a proxy is used, which is often unavoidable, but because the proxy is adopted unexamined at the moment of decision, when there is no capacity to test whether it fits the case at hand. A heuristic prepared in advance of the decision moment, and validated against an independent reference, converts an implicit substitution into a deliberate one. The condition is therefore the prior availability of a reference point, an independent estimate, a documented standard, or a comparable case already assessed for comparability, against which the proxy can be checked without extending the decision window.

The third is conditional approval. Early closure ([subsection 3.3.3](#)) fails when a decision is committed while material uncertainty remains and that uncertainty is left implicit, so that nothing in the record directs attention back to it once the assessment is procedurally complete. The protective condition is not the avoidance of closure but its explicit conditioning: a closure that names the information still outstanding, the point at which it is expected, and the circumstances under which the decision would be reopened. This preserves the operational benefit of closing within the available window while keeping the unresolved residue recoverable.

Two features of this set should be noted. First, all three are procedural or institutional in character, because all three are derived from the structure of the mechanisms, which is itself procedural. Whether conditions of a different kind also moderate the mechanisms, relational, cultural, or individual, cannot be determined deductively and is left as an empirical question for [chapter 5](#) and [chapter 6](#). Second, the protective conditions are not presented as a hierarchy or as jointly sufficient. They are the conditions the framework predicts should matter; their actual operation, interaction, and boundaries are examined empirically.

3.5. CIA completeness as an outcome variable: scope and depth

CIA completeness is the primary outcome variable of this study. It is not treated as a binary property, complete versus incomplete, but as a graded and selective outcome shaped by time pressure, uncertainty, and the distribution of information across actors. During construction, CIA rarely aims to exhaustively identify all possible impacts of a proposed change. Instead, assessments are formed under conditions that require prioritization and simplification, making variation in assessment outcomes not only likely but systematic.

CIA completeness is conceptualized along two interrelated dimensions: scope and depth. Together, these dimensions describe how the bounded-rationality strategies identified in [section 3.3](#) translate into observable limitations in what is assessed and in how far impact reasoning extends.

3.5. CIA completeness as an outcome variable: scope and depth

Scope refers to the breadth of domains, interfaces, stakeholders, and knowledge holders included in a CIA episode. Scope is constrained when relevant perspectives or impact domains are excluded from the assessment boundary. Under time pressure, decision-makers rely on heuristics, most notably formal role designation, to determine whose input is required. While this reduces coordination effort and supports timely decision-making, it systematically narrows CIA scope by privileging organizational legitimacy over epistemic relevance. Impacts that depend on cross-disciplinary coordination, informal expertise, or execution-specific knowledge are more likely to fall outside the assessment, even when consequential.

Depth refers to how far impact reasoning extends beyond immediate, first-order effects to include indirect, downstream, or longer-horizon consequences. Depth is constrained when assessment effort is truncated due to limited time or analytical capacity. Decision-makers tend to prioritize impacts that are immediately visible, well understood, or directly linked to ongoing work, local constructability, direct cost implications, or short-term schedule effects. Exploring second- and third-order consequences typically requires additional information search and cross-boundary consultation. When such effort cannot be accommodated within the available decision window, impact reasoning remains shallow, not because deeper impacts are unknowable, but because they are not explored before decision closure.

The relationship between the three mechanisms and these two dimensions is patterned rather than arbitrary. Truncated information search primarily constrains depth by limiting how far impact reasoning extends. Role-based heuristic substitution primarily constrains scope by filtering which knowledge is considered relevant. Early decision closure under unresolved uncertainty constrains both dimensions simultaneously, as residual unknowns discourage further expansion of either assessment boundaries or analytical reach. Together, these mechanisms explain why CIA, during construction, may appear formally adequate while remaining selectively narrow in practice.

Figure 3.1 represents scope and depth as orthogonal axes, producing four quadrants that describe qualitatively distinct CIA completeness outcomes. The lower-left quadrant, with a narrow scope and shallow depth, represents the highest-risk configuration, in which only local and immediately visible impacts are considered. The upper-right quadrant, broad scope combined with deep reasoning, represents the most complete assessment, though one that is rarely achievable under active construction constraints. The two intermediate quadrants reflect partial completeness: assessments that are thorough within a limited set of domains but neglect cross-disciplinary propagation, or assessments that consult a wide range of stakeholders but truncate their reasoning before exploring indirect effects. The matrix is not normative; it does not prescribe a target configuration. Rather, it provides an analytical framework for positioning and empirically comparing CIA episodes.




Figure 3.1.: CIA completeness matrix: scope and depth

Once the episodes are analyzed in chapter 5, EP-02 is positioned in the upper-right (deep-and-broad) quadrant and EP-05 in the lower-right (broad-but-shallow) quadrant Figure 5.1; the two are contrasted directly in subsection 5.6.2. The lower-left quadrant, which this matrix identifies as the highest-risk configuration, remains empty in that sample: the limitation the episodes display is not that too little was

3. Theoretical Framework

consulted, but that what was consulted was not reasoned through.

This conceptualization allows CIA completeness to function as an explanatory bridge between construction-phase decision-making and downstream manifestations such as corrective work. When CIA scope or depth is constrained, certain impacts remain latent at the moment of decision closure and become visible only when construction activities interact with unexamined dependencies or assumptions. Corrective work is therefore not treated as a direct indicator of poor performance, but as a plausible manifestation of systematically constrained assessment under bounded rationality.

3.6. Conceptual model

The theoretical framework developed in this chapter is summarized in [Figure 3.2](#). The model represents the causal chain through which construction-phase constraints translate into limitations in CIA completeness and, ultimately, into downstream corrective work. It proceeds in four steps.

Initially, the constraints of the construction phase, along with time pressures, uncertainty, and incomplete information, trigger bounded rationality because a thorough assessment cannot be completed within the limited decision timeframe. Second, decision-makers respond through three strategies: truncated information search, heuristic substitution via role-based relevance filtering, and early decision closure under unresolved uncertainty. Third, these strategies constrain CIA completeness along the two dimensions specified in [section 3.5](#): scope (breadth of domains, interfaces, and stakeholders included) and depth (how far impact reasoning extends beyond first-order effects). Fourth, when scope or depth is constrained, certain impacts remain latent at the moment of decision closure and only become visible downstream as corrective work, redesign, resequencing, or additional coordination effort.




Figure 3.2.: Conceptual model: bounded rationality and CIA completeness

EP-01 (introduced in [section 1.1](#)) illustrates the full chain shown here: a time-pressured decision window, truncated search restricted to the change's declared interface, narrow scope at the moment of closure, and a permitting interface that resurfaced as redesign five months later.

This model does not claim that bounded rationality strategies deterministically produce corrective work, nor that comprehensive CIA is always feasible or desirable under construction conditions. Rather, it proposes that the three mechanisms systematically increase the likelihood that certain impacts remain unrecognized at the moment of decision, and that this likelihood varies with the intensity of constraints across CIA episodes. This variation is precisely what the empirical chapters of this thesis seek to examine and refine.

3.7. Analytical implications for methodology

The framework informs data collection and coding by specifying what should be observable within each CIA episode. Indicators of construction-phase constraints include explicit or implicit deadlines, statements of production pressure, and evidence of incomplete or contested information at the time of assessment. Indicators of bounded rationality strategies include stopping rules ('we had enough to make a call'), reliance on dominant cues such as cost, schedule, or constructability, and decision closure despite explicitly acknowledged unresolved questions.

CIA completeness is assessed by tracing which interfaces, domains, and stakeholders were consulted (scope) and how far impact reasoning extended beyond first-order effects (depth). This tracing draws on both interview accounts and documentary evidence and is structured around the episode-reconstruction approach described in [subsection 4.4.3](#). Downstream manifestations are documented as corrective work that plausibly indicates latent impacts, redesign, resequencing, demolition, and reconstruction, or additional coordination, while avoiding deterministic causal claims about any individual episode.

Together, these analytical implications connect the theoretical framework directly to the methodological choices described in [chapter 4](#). The coding scheme ([Appendix D](#)) operationalizes each element of the framework into a set of deductive codes applied to interview transcripts and episode narratives, enabling systematic comparison across cases.

4

Methodology

This chapter explains how the research was designed and conducted. The aim is to study how Change Impact Assessment (CIA) operates in practice during construction and how decision-making under time pressure and uncertainty shapes the completeness of a CIA. The chapter follows the chain summarized in [Figure 4.1](#): research approach, design, case selection, data collection, and analysis. The study adopts a qualitative, explanatory approach. It uses a multiple-case design, with the CIA episode, not the project, as the unit of analysis. Cases were chosen to capture variation in constraint intensity, and the data come from three sources: semi-structured interviews, project documents, and a structured reconstruction of each episode. The analysis works within cases first, then across them, following the replication logic recommended for explanatory case-study research (Yin, [2018](#)).




Figure 4.1.: Research design overview

The sections that follow develop each part of the design in detail. For each methodological choice, I explain how it follows from the theoretical framework and from the research questions it is meant to answer.

4.1. Research approach

The research approach is qualitative and explanatory. The aim is not to test hypotheses or estimate effect sizes, but to explain how and why CIA narrows during construction-phase decision-making, and how those limits can leave impacts hidden until later in the project. A qualitative approach fits the topic because the phenomenon involves judgment, interpretation, and decision closure under uncertainty, processes that quantitative indicators cannot meaningfully capture on their own. The study therefore aims at analytical generalization: identifying mechanisms that recur across cases, rather than generalizing statistically to a population (Yin, [2018](#)).

4.2. Research design

The design is a qualitative multiple-case study. Each case is one construction-phase CIA episode: the period in which a proposed change is assessed, discussed, and decided while work on site continues. Focusing on episodes rather than entire projects keeps the analysis close to the decision moment where time pressure, uncertainty, and information availability interact. Several episodes are analyzed to capture variation in constraint intensity, tighter or looser decision windows, more or less uncertainty about what the change actually involves, and to strengthen the explanation through comparison across cases (Yin, 2018).

4.3. Case selection

The unit of analysis is the CIA episode; however, recruitment operates at the practitioner level. Rather than picking episodes first and then looking for the people involved, I use a critical-incident sampling logic. Experienced practitioners are invited to an interview and, at the start of the conversation, are asked to name a concrete CIA episode in which they were directly involved. The episode emerges from the interview and is then reconstructed using the interview account and the supporting project documents. This approach avoids selecting cases solely from records, which are often incomplete, and instead lets respondents identify relevant and well-remembered episodes.

To qualify, an episode must meet four essential criteria. First, it must occur during the construction phase, with site progress guiding the decision windows. Second, it should involve a change related to requirements or design, such as modifications to specifications, design details, interfaces, or execution needs, excluding purely administrative or commercial changes. Third, it must have a clear assessment and decision process, either formally documented or made visible through traceable coordination, enabling reconstruction of the assessment pathway. Fourth, it must be empirically traceable, meaning the combined interview and documentary record should allow the sequence of events and the decision-making basis to be reconstructed (Yin, 2018).

Two further sampling criteria apply across the full set of episodes. First, episodes are chosen to vary in constraint intensity, time pressure, uncertainty about the definition of change, and the amount and reliability of information available at closure, so that the framework's mechanisms can be examined under different conditions. Second, episodes are chosen to vary in downstream outcome: episodes with and without later corrective work are both included to avoid retrospective bias and to support contrast analysis. The severity of rework is not used as a selection criterion on its own. This would introduce bias by mixing the outcome with the object of analysis.

These criteria are then used in conjunction with three feasibility filters. Instead of being a collection of disparate changes, each episode must be a single, well-defined change. For respondents to remember the evaluation and for the documentation to continue to be available, it must be recent enough. Additionally, authorization to utilize the pertinent papers and availability and willingness to be interviewed are prerequisites.

The target sample is four to six CIA episodes from projects executed by the case organization in construction-intensive environments. Each episode is reconstructed from at least one interview together with a set of supporting documents. Where feasible, a second respondent involved in the same episode is interviewed to strengthen triangulation.

4.3.1. Applying the selection criteria

I applied the aforementioned standards to the case organization's register of construction-phase modifications. 339 reported changes from previous projects, ranging in value from insignificant sums to several million euros, made up the initial pool. There were three stages to the filtering process.

The mandatory inclusion criteria were used in the first step. Since rejected changes do not yield any downstream results that can be used to test CIA limits, only approved and implemented changes were retained. The documentation required to be adequate to reconstruct the assessment, impact evaluation, identified consultation, and recorded decision moment, and the changes had to occur during the construction process. Because their evaluation focuses on contractual entitlement rather than the spread of technical and execution implications that this study investigates, changes that were solely

4. Methodology

commercial or administrative (contractor claims and contract revisions) were eliminated. Together, the four retained change types, Additional Scope (Employer), Scope Change (Employer), Optimization (Contractor), and Unforeseen Situation (Employer), cover the owner-driven, contractor-driven, and site- or externally-driven origins of the change-source taxonomy presented in [section 2.2](#) ([Figure 2.1](#)). The case set gains theoretical variation as a result of this dispersion.

A value-based filter was used in the second stage to concentrate on the middle of the change population. Because they typically go through routine approval without a substantive examination, changes under about €50,000 were not included. Because they often result in distinct, protocol-heavy governance systems that lessen the time and information pressures under which bounded rationality processes become most apparent, changes above roughly €500,000 were eliminated. Despite being slightly above the €500k limit, one episode (EP-06, registered at €585k) in the final sample was kept because it was the only case with sufficient documentary depth in the Q4 cell of the sampling design (low constraint intensity combined with downstream corrective work). Therefore, the requirement to fill all four quadrants of the constraint $\times$ result matrix provides methodological justification for its inclusion. Decisions made within strict assessment timeframes are well-represented in the €50k–€500k band; unusual situations with prolonged formal review are not. Even so, the band offers a tenfold spread in stakes, sufficient to accommodate differences in constraint intensity across the final sample.

A recency filter was the third step. The primary CIA cycle was limited to changes from the previous 36 months, with up to 48 months allowed in cases where the first official approval occurred during the 36 months. This maintained the documentation's accessibility and kept the recall burden on respondents manageable.

A longlist of 32 potential episodes was generated by the three processes. Six episodes were chosen from this longlist based on the two sample criteria. The first was variation in constraint intensity, which was operationalized through references to deadline pressure in the documentation and the amount of time that passed between a request and a judgment. The second was variance in the downstream result, which was operationalized by the existence or lack of later non-conformance reports, further information requests, or updated drawings related to the modification. To make sure that every maintained type is represented at least once, whenever possible, the change-type spread mentioned above was used as a secondary constraint. [section 5.1](#) describes the resulting set.

4.3.2. Contractual context

The six episodes fall within four distinct contractual contexts of the overarching project. Contract A is the principal works contract, covering civil construction and the main physical infrastructure. Contract B is a separate works contract covering an electrical system component, with its own process timeline and interfacing parties. Contract C is a subcontracted scope nominated under Contract A. Contract D is an interface agreement with an external system operator.

Five of the six episodes (EP-01, EP-04, EP-05, and EP-06 primarily; EP-03 through its interface with Contract A) sit within or touch the Contract A context. One episode (EP-02) sits within Contract B and serves as a productive contrast case, both empirically (low time pressure, deep assessment) and contractually (a different procedural and stakeholder configuration). Two episodes raise an interface with a second contract: EP-03 between Contract C and Contract A, and EP-05 between Contract A and Contract D.

I view this contractual variation as a deliberate feature rather than a mistake or confound. Each contract has its own procedural conventions for change handling, and the cross-contract spread supports analytical generalization across procedurally distinct settings within a single organization. The episodes are not treated as fully independent observations of a single procedural mode, but rather as observations of how bounded-rationality mechanisms unfold across the contractual environments through which the overarching project is delivered. What this configuration means for the strength of the cross-case claims is discussed in [section 6.4](#).

4.3.3. A complex case

One episode (EP-04) is a complex case rather than a single change record. A site discovery triggered three linked change records, which were collectively resolved through a single combined approval by the project board on 26 October 2021. The primary record (€99k mitigation cost) is the analytical unit

for EP-04. The two linked records serve as context documents that show how the organization breaks down and recombines complex disruptions to make them procedurally manageable. That fragmentation pattern is itself analyzed as a bounded rationality phenomenon in [chapter 6](#).

4.4. Data collection

Data collection combines three sources to enable triangulation and process reconstruction: semi-structured interviews, document analysis, and CIA episode reconstruction. The interview is the entry point. The respondent identifies the episode that anchors the case at the start of the interview, and the documentary and reconstructive work that follows is built around that episode. Using three sources builds a chain of evidence, reduces reliance on any one respondent's recall, and lets me compare what was written down with what was enacted in practice.

4.4.1. Semi-structured interviews

Semi-structured interviews are the main source of empirical data. The interview design uses a two-informant strategy. It combines practitioners involved in the case with procedural experts from outside the case complex. This split lets me triangulate situated, project-specific accounts with cross-project reflections to calibrate the interpretation. Each interview focuses on five aspects of the CIA episode: (a) how the change was framed and what information was available at the outset; (b) what information was available and what was contested or missing; (c) who was consulted and on what basis; (d) how the assessment was judged sufficient; and (e) how and why the decision was closed. This focus lets the bounded rationality strategies, truncated information search, heuristic substitution, and early closure be identified without cueing the respondent to the theoretical concepts directly. Interviews are audio-recorded with informed consent and transcribed verbatim. The full interview guide is in [Appendix A](#).

Informant types. The study distinguishes two types of informants. Type A respondents are practitioners directly involved in one or more CIA episodes in the case complex. Typical roles include project manager, contract manager, system engineer, discipline lead, and site engineer. Their interviews are episode-anchored: the respondent picks the change episode that anchors the case, and the discussion then traces the CIA process for that episode.

Type B respondents are procedural experts with broad experience of construction-phase CIA across multiple projects, but no direct involvement in the case complex under study. Their interviews are cross-project reflective. They calibrate the case-specific findings against broader patterns and probe whether the mechanisms observed in the case are recognized more widely. Both types follow the same interview guide structure; the phase-two probes are adapted to the informant type.

Sampling plan. The target sample is five to eight interviews in total. Type B experts are purposively recruited based on documented experience with construction-phase CIA in comparable contractual contexts. Sampling continues until the within-case themes reach analytical saturation across episodes and the cross-project themes converge among at least two Type B respondents.

The interview sample. Six semi-structured interviews were conducted between 5 and 25 June 2026: four Type A and two Type B. The Type A interviews comprise two case-involved practitioners (INT-01, Respondent A, 5 June 2026, approximately 50 minutes, a senior contract manager in the case complex with 26 years of construction-phase CIA experience; and INT-06, Respondent F, 25 June 2026, approximately 30 minutes, a senior mechanical-engineering supervisor on the same case complex with prior international pipeline-project experience) and two adjacent-case practitioners (INT-03, Respondent C, 10 June 2026, approximately 30 minutes; and INT-04, Respondent D, 11 June 2026, approximately 50 minutes, conducted in English).

The Type B interviews comprise one procedural expert with a legal-contractual background (INT-02, Respondent B, 10 June 2026, approximately 42 minutes) and one process consultant from an external consultancy bureau (INT-05, Respondent E, 22 June 2026, approximately 20 minutes). The two case-complex interviews (INT-01 and INT-06) are complementary rather than redundant: Respondent A

reflects from a contract-management vantage point, Respondent F from a specialist-supervisor vantage point at site level. Both Type B respondents reflect cross-project from outside the case complex.

Full anonymized transcripts are provided in [Appendix B](#). The complete coding log across all six interviews, 127 coded fragments of which 96 were selected for thesis use, is provided in [Appendix C](#); the per-fragment log is available as a separate spreadsheet. Findings are reported throughout [chapter 5](#) and [chapter 6](#) with explicit evidence-weight calibration, as specified in [subsection 4.5.5](#).

4.4.2. Document analysis

Project documents are reviewed to provide context for interview accounts and to verify timelines and decision rationales. The documents include change requests, assessment records, engineering drawings, interface documents, meeting minutes, and correspondence. The document analysis identifies what was formally assessed, what uncertainty remained at decision closure, and which impacts were deferred or excluded from the formal record. It also lets me compare formal procedure with enacted practice. Cases in which the documented process and the actual assessment diverge are especially informative about the role of informal judgment and heuristic use.

4.4.3. Reconstruction of CIA episodes

Each CIA episode is reconstructed as a chronological narrative covering five stages: the change proposal, the assessment process, decision closure, implementation, and, where applicable, downstream corrective work. [Figure 4.2](#) shows this structure with the evidence sources and analytical focus for each stage.




Figure 4.2.: Structure of a reconstructed CIA episode: stages, evidence sources, and analytical focus

The change-proposal stage draws mainly on documentary evidence and interview accounts to establish the origin of the change, its urgency, and the initial state of information. The assessment-process stage is the analytical core of the study. It is examined primarily through interview data to identify who was consulted, what information was available or contested, and which bounded-rationality strategies shaped the assessment. Decision closure is examined through both documents and interviews to establish what residual uncertainty remained at the point of commitment. The implementation and corrective-work stages are examined to see whether and how latent impacts, those not recognized at closure, later became visible as corrective effort.

The reconstruction brings together interview and documentary evidence to identify critical decision points, information gaps, and moments where impacts were recognized, deferred, or overlooked. This narrative structure is the foundation for within-case process analysis. It also enables cross-case comparison by ensuring the same stages are examined consistently across episodes.

4.5. Data analysis

Data analysis follows a theory-informed thematic and process-tracing approach in three steps.

4.5.1. Step 1: within-case analysis

Each episode is coded against a deductive scheme drawn directly from the bounded rationality framework: (1) construction-phase constraints, time pressure, uncertainty, incomplete or contested information; (2) bounded rationality strategies, truncated search, heuristic substitution, early closure; and (3) CIA completeness outcomes, scope and depth. Inductive codes are added where recurring patterns emerge that the deductive scheme does not capture. The full coding scheme is in [Appendix D](#).

4.5.2. Step 2: process tracing

The reconstructed episode narratives are analyzed as process data. The aim is to see how constraints and strategies unfold over time and how decisions are reached. Process-oriented analysis supports mechanism inference by linking observed sequences, for example, escalating schedule pressure leading to a narrowing consultation and then early closure, to outcomes such as limited scope or depth (Langley, 1999). This temporal perspective is crucial for telling apart a mechanism's activation from a mere coincidental occurrence.

4.5.3. Step 3: cross-case comparison

Recurring configurations are identified across episodes. The explanation is strengthened through replication logic and through systematic consideration of plausible rival explanations (Yin, 2018). Where downstream corrective work occurred, it is analyzed as a potential manifestation of latent impacts, while avoiding deterministic causal claims that the analysis cannot support.

4.5.4. Quote selection

Quotes in [chapter 5](#) are selected on three explicit criteria. The first is representativeness of a coded pattern: the quote illustrates a recurring observation across multiple sources or cases rather than an isolated remark. Where several quotes in the coding log support the same pattern, the one presented is the one whose phrasing most directly captures the mechanism. The others remain available in the coding log for verification. The second is clarity in capturing the mechanism: the wording explicitly shows the bounded rationality strategy at work, so the reader can verify the code–quote relationship directly without much interpretive scaffolding. Where the most compelling quote is interpretively richer but less explicit, two quotes are presented, a representative one and a striking one, and the distinction is noted. The third is traceability: each quote corresponds to a specific document identifier in the Document Log and a code reference number in the Coding Log, enabling the reader to locate the source and independently verify the coding decision. Where contrasting quotes within the same case exist, for example, substantive versus shallow disciplinary inputs in the broad-but-shallow consultation pattern (EP-04), both are presented to avoid selectively representing the within-case variation. Quotes are not chosen to support a predetermined narrative; they are chosen to illustrate the patterns identified through systematic coding, with the coding log providing the complete evidence base.

4.5.5. Mechanism evidence weighting

The characterization of bounded rationality mechanisms per episode (Yes, Partial, or No in [Table 5.3](#)) is supplemented by an evidence-weight indicator. The indicator reflects both the number of coded observations supporting that mechanism within the case and the diversity of source types from which those observations come. Three categories are used. Strong evidence means three or more coded observations across at least two distinct source types, for example, a core team analysis memo plus individual disciplinary correspondence. Moderate evidence means one or two coded observations, or several observations from a single source type. Preliminary evidence means a single observation from a single source. Where a mechanism is coded as absent or partial, the indicator describes the basis for that classification, for example, positive evidence of a phased approach as in EP-02, rather than simple absence of observation.

Two calibration rules operate, corresponding to the two kinds of claim made in this study. The first applies to mechanism coding within an episode, where the evidence base is documentary and the

4. Methodology

rule is as stated above. The second applies to cross-case findings that are visible only in the interview material, including conditions and patterns that the documentary record cannot capture by construction. For these, Strong denotes convergence across three or more interviews from at least two different vantage points relative to the case complex, inside it, adjacent to it, or outside it; Moderate denotes two independent interviews from different vantage points; and Preliminary denotes a single source. Where a finding is supported by both documentary and interview evidence, the two rules are applied jointly and the finding is reported at the lower of the two weights. This distinction matters because several of the findings developed in [chapter 6](#) concern informal practice that is, by its nature, absent from the documentary record; applying the documentary rule to them would understate evidence that is nonetheless multi-source.

The weighting is not used to decide the code itself: that remains based on whether the mechanism is substantively present or absent. The weighting is used to calibrate the strength of cross-case claims in [section 6.1](#) and the confidence with which patterns are described in [section 5.4](#), for example, ‘strongly observed in EP-01 and EP-04, with preliminary evidence in EP-06’ rather than the flatter ‘observed in four of six cases’. The evidence-weight indicator per case and mechanism combination is recorded in the Coding Log and summarized in an extended version of [Table 5.3](#).

4.6. Quality criteria and validity

Three things support the research quality: triangulation across sources, transparent documentation of the analytical steps, and systematic episode reconstruction. The study keeps a clear chain of evidence from raw data through case narratives to cross-case claims. A structured interview guide ([Appendix A](#)) and a coding scheme ([Appendix D](#)) keep the analysis consistent across episodes.

To support credibility, the analysis explicitly aligns evidence, codes, and claims with the study’s research questions and theoretical framework. Member checking, sharing episode reconstructions with participants for accuracy verification, is used where feasible to reduce recall bias and interpretive error. Quality is further supported by reporting sufficient contextual detail to enable transferability and by explicitly reflecting on interpretive choices and potential bias. The approach is consistent with widely cited criteria for excellent qualitative research, including credibility, rich rigor, and meaningful coherence (Tracy, 2010).

4.7. Ethical considerations

The study follows established ethical principles for qualitative research, with particular attention to informed consent, confidentiality, and the responsible handling of sensitive project information. Informed consent was obtained before data collection. Participants were informed of the purpose of the study, the voluntary nature of participation, and their right to decline any question or withdraw at any point without consequence; interviews were audio-recorded only with explicit permission, as documented in [Appendix A](#). Access to recordings, transcripts and documents is restricted to the researcher and supervisory team, and all materials are stored in line with TU Delft data management requirements.

Anonymization operated in two layers. The change records were provided by the case organization in already anonymized form: the researcher had no access to personally identifiable data or unfiltered client identifiers and did not perform the initial anonymization. In those records, partner organizations appear as generic placeholders and individuals are referenced by role. A second layer was applied to the material reproduced in this thesis, in which the case organization is referred to as such rather than by name and internal change-record numbers are retained only as opaque case markers. This regime was a precondition of data access and is consistent with the approval obtained from the TU Delft Human Research Ethics Committee.

The analysis is conducted reflexively and non-evaluatively. Explanations are framed in terms of decision constraints and mechanisms rather than individual error, and this framing was stated in the interview preamble so that respondents understood the purpose of the research before giving their accounts.

5

Findings

This chapter presents the results of the study, based on an analysis of six CIA episodes. Rather than discussing each episode separately, the analysis is organized by theme, following the structure of the theoretical framework. It first examines how CIA operates in practice, then identifies the limitations observed in each episode and the strategies used to address them. Next, it assesses the completeness of CIA in terms of scope and depth. Finally, it analyzes downstream effects, focusing on corrective actions taken after implementation. Patterns that appear across multiple episodes are highlighted in each section and further developed in [chapter 6](#).

5.1. Overview of CIA episodes

Before the cross-cutting analysis, it is useful to be precise about what constitutes an episode. An episode is a single construction-phase change, reconstructed as a bounded decision story: it begins with the initial identification or suggestion of the change, the progresses through evaluation and formal consent (the initial and subsequent approvals), and concludes with execution and any subsequent rectifications. The project serves as the backdrop; the episode represents the trajectory of a singular transformation inside it, analyzed according to the five-stage framework outlined in [subsection 4.4.3](#). The six episodes were selected from a longlist of 32 candidates using the procedure outlined in [section 4.3](#). They were intentionally selected to encompass four categories of change: owner-initiated scope, contractor optimization, unexpected site conditions, and cross-contract interface modifications, allowing the framework to be evaluated across many decision types instead of being confined to a single category. [Table 5.1](#) provides a comparative summary; the subsequent profiles present each occurrence sequentially.

The first additional watercourse crossing on the major works contract (Contract A) was EP-01, where a dam had to be demolished and a culvert installed in its place. Twenty-three days separated the first dated document from the first approval since a defined date for external notice had to be approved. The examination was strict yet procedurally correct and the crossing itself was adequately analyzed. It did not catch that the authority responsible for the watercourse had terms of its own which were slightly outside the legal assessment border. They came after the first approval and the design work had to be redone for them. It is the story of an assessment that answered its own question well, but not the question that counted afterward.

EP-02 is the case that stands out. A part of the export cable route that went by a high-voltage substation was rerouted, and more conduit was added. This happened on Contract B, which was the only episode in the sample that involved a separate contract. There were 321 days between the first dated document and the first approval, an unusually generous window. The core team used that room to look at both the change's instant effects and its history over time. They also asked what would happen if this connection ever had to be shut down. Because of that question, it was decided to build a reserve element next to the working one, weighing the cost of doing so now against the practical risk that would happen over the next 25 years. No corrective work followed. This is the story of what a thorough assessment looks like when time is not the binding constraint.

EP-03 is located where a subcontracted scope and the primary works (Contract C - A) meet. Thirteen days separated the first dated document from the first approval, and at the time of the decision, a crucial input, the partner party's terms and price, was still pending. The team closed anyway. What made the

5. Findings

closure operationally safe was not the timing but how the approval was written: it was made conditional on the partner terms arriving, with a clear "re-open" clause in case they came back differently. Instead of being activated when those terms came, the condition vanished. No rework and a change of €85k. Because it was written to be reopened, the story of an early closure did not turn into an early mistake.

In the beginning of EP-04, a protected species was found on the site, which is illegal to remove because of nature conservation laws. At the same time, one find led to three linked changes: a way to mitigate the problem, a dispensation route, and a consequence for site access. Rather than handle them separately, the core team bundled them into a single board decision. The consultation table was wide; nine different roles contributed to the assessment memo. But the amount of work that each part did was not equal, and some of the contributions were more like acknowledgments than real analyzes. By the time of the second approval, a multi-month access delay had turned into a €99k schedule-recovery premium. The delay was bought out in the end instead of being planned for. The story is about a table that was very wide but not very deep.

EP-05 is the story in the sample that costs the most. Two structures had to be moved at the boundary with an external system operator (Contract A–D), and they had to be done in the shortest amount of time possible because of a production deadline for a long-lead cable. The core team closed the decision through a combined approval that merged the first and second approval steps into one moment. The milestone was then safeguarded. A unfavourable municipal decision a few months later overturned the technical basis on which the shutdown had been based, so things had to go back to how they were before. Then there were about €190k in mobilization costs, and some of the civil works that were already there were torn down. That's the story of an unanswered doubt that the moment of decision could not settle, and that resurfaced downstream at full cost.

The biggest modification in the sample is EP-06, valued around €585k, and took 617 days from first documentation to a signed change order. The change modified the site access roads and bridge ramps that carried heavy transport onto the site. What made it analytically interesting was not the size but the shape of its time-pressure profile. Most of those 617 days there was no acute pressure at all. The assessment covered predicted transport and execution requirements. What was never tested was the use phase, and notably the agricultural users that would have to cross those same entrances once the site was handed over. Fourteen road locations were then lowered and a temporary transport route created at a cost of roughly €218k. It's the story of a blind spot that time pressure cannot explain by time constraints.

Against this overview, the following sections analyze how CIA was enacted across the episodes ([section 5.2](#)), the constraints under which it took place ([section 5.3](#)), the bounded-rationality strategies observed ([section 5.4](#)), the resulting completeness in scope and depth ([section 5.5](#)), and the downstream corrective work ([section 5.6](#)).

Table 5.1.: Overview of CIA episodes

	Contract	Change source	Change description	Window (days)	CIA quadrant	Corrective work	Key BR strategy
EP-01	A	Additional scope (Employer)	Scope addition; external permitting interface left outside initial scope	291	Broad-but-shallow	Yes	B2 + B3
EP-02	B	Additional scope (Employer)	Employer-added scope on an electrical-system component, assessed across its service life	513	Phased / deep	No	Contrast case
EP-03	C ↔ A	Scope change (Employer)	Scope change at the subcontract–main-works interface; partner terms unknown at decision	54	Compressed, conditioned	No	B3 with conditional pause
EP-04	A	Unforeseen situation (Employer)	Unforeseen site condition triggering three linked changes, resolved in one combined decision	50	Broad-but-shallow	Uncertain	B1 + B2 + B3
EP-05	A ↔ D	Risk management (Employer)	Interface change with an external system operator under a hard production deadline	96	Compressed (combined approval)	Yes	B1 + B2 + B3
EP-06	A	Additional scope (Employer)	Large access and scope change; use-phase interface with surrounding land users not assessed	617	Broad-but-shallow over time	Yes	B1 + B3

5.2. CIA as conducted in practice

This section answers SQ1 by describing how CIA was enacted across the episodes: the steps taken, the actors involved, the artifacts used, and the points at which decisions were formalized. Instead of walking through each episode in full sequence, I first set out the common process structure and then highlight the variations that matter analytically.

5.2.1. Common process structure

Across all six episodes, CIA followed a common four-stage procedural structure: initiation, assessment, formal authorization, and implementation. Initiation happened either externally (a contractor-submitted change request) or internally (an early warning raised by the Change Manager). In all cases, a dated document marked the formal start of the assessment process. The assessment phase combined two document streams that appeared in every episode: the contractor's price proposal with scope description, and the core team analysis memo, a project team synthesis prepared for the steering group decision. Formal authorization followed a two-step pattern: the first approval gave in-principle approval to continue the change, and the second approval gave definitive financial and contractual approval. Implementation was finalized by a signed change order, which closed the procedural cycle and committed the parties to execution. This four-stage structure was identifiable across all four contractual environments (Contracts A, B, C, and D), with the steering group functioning as the decision authority across all six cases.

Disciplinary inputs varied in number across episodes but were consistent in their procedural placement. Inputs from the Contract Manager (on contractual implications), the Financial Manager (on cost consequences), and, where the change touched the external environment, the Environmental Manager (on permits and stakeholder communication) appeared in five of six episodes. Health and safety and project secretary inputs appeared in the more complex cases (EP-04 and EP-05). Process coordinator inputs appeared consistently across cases within Contract A. These disciplinary contributions were typically captured as short email or memo entries and consolidated into the core team analysis before reaching the steering group. The formal decision point was identifiable in all six episodes through a documented steering group approval, with the second approval representing the binding moment of resource commitment. Decision authority sat with the steering group throughout; the Change Manager served as procedural facilitator.

5.2.2. Variations that matter analytically

There are three analytically important variations across the episodes. These are the phased versus compressed decision architecture (related to B3), the extent of disciplinary consultation (related to B1) and the total process duration (related to external time-pressure conditions).

Variation 1: phased vs compressed decision architecture. The most analytically important distinction was the temporal distance between the first and the second approval, and this distinction partly patterns with the contractual context. The only episode in Contract B, EP-02, showed the fully phased pattern intended by the procedural design: first approval on 25 May 2021, second approval on 1 December 2021, with 191 days of iterative refinement in between. Technical and contractual items left open at the first approval were addressed in three rounds of contractor offers and Employer comments. Similarly, EP-01 and EP-04 within Contract A also exhibited a phased logic (but with shorter refinement intervals: twelve months and approximately five months, respectively). The other 3 episodes had a compressed process with the first and second approval being merged or nearly merged: EP-05 had a merged approval on 13 September 2021, EP-06 both approvals on 31 August 2021, and EP-03 a combined consultation round within 20 days. This compression of the two approval moments should be differentiated from the bundling of several change records into one decision that took place in EP-04 (section 4.3.3) while its two approval moments were still five months apart. They are two different forms of procedural compression, only the first of which is relevant here. The compressed pattern is therefore mostly found in Contract A and in cases of cross-contract interface (EP-03, EP-05). The pattern of phase deliberation is most evident in Contract B. This pattern indicates that the speed of the CIA

process is affected not only by time pressure but also by the contractual environment. This point is developed in [section 6.1](#) as evidence that B3 (early closure) is an institutionally available procedural option, not just a cognitive shortcut.

Variation 2: degree of disciplinary consultation. A second variation was in the number of disciplinary roles that were embedded in the dossier (from about four roles (EP-02) to nine roles (EP-04)). This is important in two ways. First, the scope of formal consultation is not coextensive with the scope of substantive review. As [section 5.4](#) describes, the broad consultation by EP-04 was concurrent with multiple disciplinary contributions with little substantive content. The broad-but-shallow pattern of consultation was strong in EP-01 and EP-04. Second, breadth of discipline relates imperfectly to the procedural compression discussed above: EP-05 involved a high disciplinary input under acute external time pressure with combined-approval compression; EP-03 involved a narrower internal consultation with a strict combined-approval cadence; EP-01 involved a mid-range disciplinary footprint with phased refinement. This variation is not a measure of quality but evidence of two partially independent dimensions of CIA practice. There is also no strong stratification by contractual environment: broad-but-shallow patterns appear within Contract A, and would, on the present sample, plausibly appear in Contract B too, although this remains testable.

Variation 3: process duration and external triggers. The third variation was the total time of the process from the first dated document to the signed change order. It varied between approximately three months (EP-05) to more than two years (EP-06). The drivers respect the constraint conditions of [section 5.3](#) and the contractual environment. Short overall durations were either due to strong external time pressure (EP-05, with a deadline from the external system operator under the Contract D interface) or quick contractual cycles (EP-03, where the contractor's pre-positioned work generated urgency for resolution under Contract C). Long total durations were associated with either intentional phased refinement within Contract B (EP-02) or unresolved downstream issues that created multi-year iteration cycles within Contract A (EP-06, where the road-height correction generated documents over a period of more than two years). EP-04 was in the center. This variation is important because it shows that procedural duration is not a uniform proxy for any one underlying condition. The time-pressure dimension developed in [section 5.3](#) is expressed at the level of the assessment window, between first proposal and first approval, and not the entire process cycle, and this is true regardless of contractual environment.

5.3. Constraints observed

This section answers SQ2 by characterizing the time pressure, information incompleteness, and uncertainty that shaped each CIA episode. [Table 5.2](#) gives a comparative intensity rating across all six episodes; the prose that follows works out the most analytically significant constraint configurations.

5.3.1. Time pressure

Time pressure operates across the six episodes in two analytically distinct forms. The first is operational time pressure within a single CIA cycle: the interval between when a change is identified and when a decision must be made to keep the construction sequence moving. The second is portfolio-level time pressure: the rate at which new change requests enter the assessment pipeline relative to the team's capacity to assess them. The episodes show both forms, but in different combinations.

EP-05 shows the most acute operational time pressure, with a cascade of near-simultaneous external deadlines culminating in a cable production deadline with no slack. EP-03 shows a comparable operational profile in compressed form, with a 54-day window from the first dated document to the signed change order, driven by pre-positioned works at the partner interface. EP-04 sits in the middle: a nature-conservation dispensation timeline operating in months rather than days, but with a hard biological-window deadline that could not be negotiated. EP-01 and EP-02 show low operational time pressure, but for different reasons. EP-01 because the external regulator's four-week notification requirement allowed, and forced, a compressed but procedurally orderly assessment; EP-02 because

5. Findings

Table 5.2.: Constraint intensity by episode

	Time pressure (A1)	Info. incompleteness (A2)	Uncertainty (A3)	Dominant constraint	Key evidence (Coding Log refs)
EP-01	Low	High	Med	A2 (info)	Refs A-001, A-002; late notification of change, four-week external deadline
EP-02	Low	Med	Med	A2 (info)	Refs A-006–008; new soil data arrived in time; permits open but resolvable
EP-03	High	High	High	A2 (info)	Refs A-011–013; pre-positioned works create contractual urgency; partner pricing unknown
EP-04	Med	High	High	A3 (uncertainty)	Refs A-015–018; nature-conservation dispensation pending; cost a black box
EP-05	High	High	High	A3 (uncertainty)	Refs A-023–025; cascade of hard deadlines including the cable deadline; municipal ruling open until day –1
EP-06	Med	High	High	A2 (info)	Refs A-029–031; no requirements set for top and foot arc of road entry; social unrest escalates over time

the contractual milestone structure of Contract B allowed iterative refinement across two approvals six months apart.

EP-06 is analytically the most interesting case in this section: the total process duration is 617 days, but time pressure is rated Medium because it intensified non-uniformly across the cycle, becoming acute only when social unrest among landowners surfaced and demanded a fast resolution.

Three interviews corroborate the episode-level time-pressure pattern from outside the single-case documentary record. Respondent A, drawing on 26 years of practice across multiple projects, names it as the single most consistent constraint he has encountered:

Tijdsdruk ... weegt het zwaarst. Als ik zo terugkijk naar de afgelopen 26 jaar, is tijdsdruk op elk project geweest. Dan heeft het op elk project veel invloed gehad. Absoluut.

[Respondent A, CL-027]

Respondent B, a procedural expert calibrating across many projects rather than this case complex alone, identifies time pressure as what separates a well-handled change from a poorly handled one (CL-112). Respondent C describes the same condition in a market-driven deadline context rather than a regulator-driven one:

Het is al verkocht productie en dat moet af zijn in juni. Dus ja, daar zit je al met een tijdsdruk. En ja, je doet bij elke wijziging die binnenkomt is de aanneming 'het is nodig'.

[Respondent C, CL-206]

Document analysis and three converging interviews, from inside the case complex, from a cross-project procedural role, and from an adjacent project, agree that time pressure functions as the precipitating condition rather than one of three co-equal triggers. This finding is treated as Strong and carried forward into the interpretation in [subsection 6.1.1](#).

5.3.2. Information incompleteness

Information incompleteness at the moment of decision is the most consistently rated constraint across the sample, scored High in five of six cases ([Table 5.2](#)). The incompleteness takes three forms.

The first is incomplete cost information. In EP-01, EP-03, EP-04, EP-05, and EP-06, the contractor's final price proposal was either absent or only available as an order-of-magnitude estimate at the moment the first approval was taken. In EP-04 the assessment memo recorded that the consequences were still to be quantified by the contractor.

The second is incomplete interface information. In EP-03, the partner organization's pricing and terms were unknown. In EP-06, essential technical requirements for the use phase, the top arc, foot arc, and straight-section dimensions of the road entries, had not been set at all.

The third is incomplete regulatory information. In EP-05, the municipal ruling on the permit was unknown until one day before the binding decision was due. EP-02 is again the contrasting case: the relevant soil investigation data arrived in time to be analytically integrated before the first approval, and the permits were known to be resolvable, so information incompleteness operated at the Medium rather than the High level.

5.3.3. Uncertainty

Uncertainty is distinct from information incompleteness in the framework ([section 3.2](#)). Incompleteness concerns information that exists but is not yet available. Uncertainty concerns information that does not yet exist because the underlying situation has not yet stabilized. Three of the episodes (EP-03, EP-04, EP-05) show High uncertainty driven by interface dynamics that were still unfolding at the moment of decision.

EP-04 combined regulatory uncertainty, whether the nature-conservation dispensation would be granted, with biological uncertainty about the protected species' behavior over the construction window. EP-05 combined regulatory uncertainty (the municipal ruling) with social uncertainty about whether the affected community would accept the alternative siting. EP-06 shows a slower-moving form of uncertainty in the relational and juridical track with landowners, potential claims for unusable headlands and

5. Findings

damaged agricultural plots that emerged over months rather than days. The analytical implication, developed in [subsection 6.1.1](#), is that uncertainty more often co-occurs with, and amplifies, time pressure than it functions as an independent, dominant constraint. Only in EP-04 and EP-05 is uncertainty plausibly the dominant constraint in its own right.

5.3.4. Upstream conditions: contracting-stage origins

The three constraints documented above ([subsection 5.3.1](#)–[subsection 5.3.3](#)) are not fully explained by conditions internal to the CIA episode itself. Two interviews point to conditions established earlier, at the contracting stage, that shape the severity of the constraints the construction-phase CIA team subsequently inherits.

The first is a design-maturity pathway. Respondent A describes the case complex's Front-End Engineering Design (FEED) package as known to be incomplete at the point construction began, with the decision to proceed taken deliberately, in full awareness of the rework consequences that would follow:

Antwoord was eigenlijk: nee, het is niet goed genoeg. Later hoorde ik bevestigd worden dat de ontwerpbureaus X en Y hadden gezegd: jongens, wat doen jullie nu? Wij zijn nog aan het engineer. Het is inderdaad ook niet klaar.

[Respondent A, CL-003]

Respondent D, working on a different, currently pre-execution project, describes the identical condition unfolding in real time rather than in retrospect:

We really have to get the engineering locked down. That's the most important. At the moment, it's just not locked down enough.

[Respondent D, CL-312]

The second is a procurement-fragmentation pathway. Respondent C traces it to Final Investment Decision (FID) price pressure: when the FID requires a lower contract price than a single-party offer allows, the project is split into more, smaller contract packages, which increases the number of interfaces and, in turn, the number of changes that must each be separately assessed under the same time constraints:

Dat heeft te maken met FID. Ze hadden één partij die het had aangeboden voor X bedrag, en dat was te duur. Toen hebben ze het opgesplitst in kleine stukjes om lagere prijs te krijgen. (...) Maar ja, je managementorganisatie wordt veel groter.

[Respondent C, CL-221]

Respondent F corroborates the procurement-fragmentation pathway with a concrete case-complex instance, a scope split roughly 300 metres from a compressor station, and identifies its downstream effect as a specific communication-suppression pattern at the resulting interface, rather than simply more interfaces:

Normaal gesproken: de pijpleiding-aannemer bouwt heel de pijpleiding en communicatie tussen die twee partijen is wat het is. Maar nu zit daar een aannemer tussen die alle kaarten tegen de borst aanhoudt als het gaat om allerlei beslissingen. Het wordt heel erg lastig om daar goede besluiten en goede communicatie te onderhouden.

[Respondent F, CL-509]

The design-maturity pathway is supported by case-complex documentary context and INT-04 evidence (Strong). The procurement-fragmentation pathway is supported by two converging interviews from different vantage points, INT-03 (adjacent case) and INT-06 (case complex) (Moderate). Both pathways describe the same underlying pattern: a compromise that is individually rational at the contracting stage produces the constraint environment within which the per-change mechanisms documented in [section 5.4](#) subsequently operate.

5.4. Bounded rationality strategies observed

This section answers SQ3 by identifying which bounded rationality strategies were used across the episodes and the conditions under which each occurred. Across the six episodes, all three bounded-rationality strategies in the framework appear. Truncated information search (B1) is observed in four of six episodes, with strong evidence in EP-01 and EP-04, preliminary evidence in EP-06, and a positive contrast case in EP-02. Heuristic substitution (B2) appears in four of six episodes, with preliminary evidence where coded. Early closure (B3) is observed in five of six episodes, with moderate evidence in EP-05 (two coded observations covering both the pre-emptive offer request and the combined approval route) and preliminary evidence elsewhere. EP-02 is the only episode in which B3 is positively absent. The varying evidence weight across mechanism and case combinations matters for the cross-case claims developed in [section 6.1](#): refinements of the framework are presented at weights that reflect both the binary code assignment and the evidence underpinning it.

5.4.1. Truncated information search (B1)

B1 manifests in the case episodes as the broad-but-shallow consultation pattern: many disciplinary roles are formally consulted, but the substantive content of their inputs is thin. Respondent E, a procedural-process consultant working across projects, names this pattern self-aware:

Ik heb projecten meegemaakt waarin het proces zo gestructureerd is dat al die check-poortjes zijn ingebouwd. Dat werkt ook niet, want dan krijgt iedereen alles te zien. Heel veel mensen vinden van heel veel niks, dus is het 'geen opmerkingen' en door. Heb je hun buy-in? Of is dat meer een papieren werkelijkheid en vertraagt het proces ook enorm.

[Respondent E, CL-408]

Respondent F describes a concrete case-complex instance of the same mechanism, a small specification change that appeared low-impact at decision closure but produced a large downstream cost cascade once its implications were traced:

Er was van tevoren bedacht dat we 10-micronfilters zouden plaatsen in een compressorstation. Nu wordt dat één micron. Op zich nog niet zo'n gekke wijziging, die kan je daar gewoon doorheen zetten. Maar de gevolgen daarvan is wel dat de leiding nu op een andere manier gereinigd moet worden. (...) Dan praat je opeens over een serie compressoren met droge units die weggezet worden tegen pak hem beet ergens 30.000 per dag.

[Respondent F, CL-503]

Respondent F's procedural diagnosis names the cognitive closure that the framework's B3a mechanism identifies: attention moves to the next decision once the present one is procedurally complete, regardless of whether the downstream depth has been explored:

Dan is er waarschijnlijk de focus op een gegeven moment gelegd naar 'nou, we hebben het afgetikt, dus we weten wat we moeten doen' en dan wordt de focus naar het volgende onderwerp gelegd.

[Respondent F, CL-504]

EP-01 and EP-04 show this most strongly. In EP-04, nine roles contributed to the dossier, but several inputs consisted of short boilerplate acknowledgements, satisfying procedural breadth without contributing analytic depth. EP-06 shows a different B1 sub-type: discipline silos in which the contract-engineering and environment teams each addressed only the change-phase impact, with neither addressing the use-phase agricultural interface that later drove the downstream rework. EP-02 is the positive contrast: the iterative refinement structure under Contract B provided the time and the process timing to deepen rather than broaden the consultation, producing strong evidence of B1 absence.

The interviews add why B1 occurs, beyond what the episode documents show on their own. Respondent A locates the cause not in oversight but in the relational economics of collaboration: searching too deeply, particularly on cost, signals distrust and can damage the working relationship the project depends on:

5. Findings

Maar als je [accountantcontrole] gaat doen, dan zeg je eigenlijk ook: ik vertrouw jou niet. En als je in dit soort projecten tegen elkaar zegt 'ik vertrouw jou niet', dan heb je ook een ander probleem.

[Respondent A, CL-011]

Respondent C provides an unusual second corroboration: he describes himself, not the contractor, displaying B1. Equipment rentals that should have been scoped upfront were instead released as individual site instructions, with costs accumulating until year-end:

Wat bij mij in mijn contract fout was gegaan, is alle verhuurmateriële. Ik had niet van tevoren bedacht wat we nodig gingen hebben. Dat is dus in losse site instructies eruit gegaan, waardoor de kosten zich opstapelden tot het einde van het jaar.

[Respondent C, CL-214]

Self-acknowledged B1 is analytically valuable because interview respondents typically defend past choices rather than volunteer their own lapses. It strengthens the reading that B1 is a structural condition available to any actor facing a binding time-versus-information trade-off, not an error mode specific to one party. After documentary evidence and three converging interviews, B1 is treated as Strong.

5.4.2. Heuristic substitution (B2)

B2 appears in three documented subtypes across the episodes. The first is standard-pattern substitution: EP-04 applied a risk discount on nature-conservation dispensation costs, and EP-05 reduced a high-dimensional siting decision to a two-by-two scenario matrix. The second is schedule-driving substitution: at EP-04's second approval, the €99k acceleration heuristic substituted the question of what the CIA impact was with the question of whether the change fitted the milestone envelope. The third is a candidate sub-type surfaced by INT-01, pick-your-battles substitution (CL-006), in which low-impact changes are deliberately closed on relational grounds. The third sub-type rests on single-interview evidence and is reported as Preliminary, pending further corroboration.

INT-06 introduces a fourth candidate sub-type, confidence-substitution. Where the previous subtypes substitute content (a prior case, a milestone-fit question, a relational closure, an internal narrative, a lowest-risk option), confidence-substitution substitutes form: the actor accepts information not because it is verified but because it is presented persuasively. Respondent F describes the pattern directly:

Mensen maken keuzes die eigenlijk die keuze niet hebben kunnen maken of onvoldoende geïnformeerd zijn. (...) Dat kan ook omdat ze denken dat ze de juiste informatie hebben omdat degene die die informatie verstrekt heeft zo overtuigend is dat hij dat gewoon goed kan brengen.

[Respondent F, CL-506]

Analytically, confidence-substitution differs from the other sub-types in what it substitutes for. The decision-maker substitutes verification with presenter confidence. The sub-type rests on a single source and is reported as Preliminary; if corroborated, it would suggest that B2 has both content-substitutions (the existing sub-types) and form-substitutions (this one).

Two further sub-types, each from a different respondent, extend the typology. Respondent B names an internal-narrative sub-type: actors selectively frame the elements they have won in a negotiation, downplaying what they have given away, so the result reads as a success when reported internally (CL-107). Respondent D names a risk-minimization sub-type: rather than substituting the best option, the actor substitutes the lowest-risk option and defends the choice as rational under uncertainty:

You're making a decision around risk. (...) I know this change option B has got the lowest risk. It's not the best option, but it's got the lowest risk.

[Respondent D, CL-319]

With these two additions, B2 now has four interview-sourced sub-types (pick-your-battles, internal-narrative, risk-minimization, confidence-substitution) alongside the two documentary sub-types (standard-pattern, schedule-driving). All four interview sub-types remain Preliminary pending further corroboration.

5.4.3. Early decision closure (B3)

B3 appears in five of six episodes, with the combined approval pattern merging the first and second approval into a single closure moment as the dominant institutional vehicle. This pattern is stratified by contractual context: combined approval closures are concentrated in Contract A and at cross-contract interfaces (EP-03, EP-05, EP-06), and are absent in the single Contract B episode (EP-02). EP-03 shows that B3 can be operationally safe when paired with an explicit waiting condition for later information, in this case the partner's pricing and terms, supporting the conditional approval protective condition introduced in [section 3.4](#). EP-05 shows B3 in its most consequential form: the combined approval taken under the cable deadline protected the milestone but exposed the project to the full downstream cost of the negative municipal ruling. The interviews add two further refinements to B3, developed below: a second B3a vehicle distinct from the combined-approval pattern, and a second scale at which early closure operates entirely.

The interviews add a second B3a vehicle that the case-complex documents do not make explicit. In the site-instruction route, execution precedes both CIA and formal closure: work proceeds first, and the multi-signature approval process ends up ratifying what has already been done rather than approving it in advance. Respondent C describes the pattern directly:

Ja, vooral omdat werkzaamheden nu uitgevoerd moeten worden. Dingen die we niet voorzien hadden. En dan zeggen we tegen de aannemer: 'ga dit nu uitvoeren, want we hebben het nu nodig', maar dan is het al uitgevoerd voordat we het beoordeeld hebben.

[Respondent C, CL-203]

Respondent D adds a normative distinction the case-complex documents do not articulate: where the site instruction is safety-driven, work has to stop or proceed differently because the alternative is unacceptable, and execution-before-CIA is legitimate; where it is substance-driven, procedural escalation is being avoided because it is inconvenient, and execution-before-CIA is the closure shortcut the framework should track. Respondent E adds the governance form this distinction implies in practice:

Daar moet je goede procesafspraken over maken hoe site instructies in elkaar zitten. Wie mag dan die instructies schrijven, afhankelijk van de impact? Daar kan je breakgrenzen inbouwen. Bijvoorbeeld: de supervisor die buiten loopt, mag tot 10K instructies schrijven.

[Respondent E, CL-415]

After three converging instances (INT-03, INT-04, INT-05), the site-instruction-backfill sub-type is treated as Strong, alongside the combined-approval pattern as the two documented B3a vehicles.

A second, portfolio-level scale of early closure (B3b) is also visible only in the interviews. Respondent A describes a periodic bundling of many individual change decisions into a single negotiated settlement, on the ground that the carrying overhead of separate handling exceeds the marginal informational gain of doing them one by one:

Vorig jaar hebben we een settlement afgesloten ... we spreken een nieuwe oplever-datum af, maar we werken ook met een clean slate, dus alle wijzigingen voor deze datum bestaan niet meer. En we rekenen daar één bedrag voor af.

[Respondent A, CL-024]

Respondent B corroborates this independently, from outside the case complex, describing forty individual changes bundled into a single negotiation from his own contract-management practice (CL-105, CL-106), and adds a refinement the first account does not make explicit: below a tolerance threshold of roughly 10% of contract value, portfolio-level closure occurs not through deliberate settlement but through routine non-attention (CL-111). After two independent interviews from different vantage points, the B3a and B3b distinction is treated as Moderate.

5.4.4. Protective conditions

This section answers SQ5 by identifying the conditions under which the strategies documented above did not translate into reduced CIA completeness. Where [section 5.3](#) and [section 5.4](#) document the constraints and mechanisms that narrow CIA, the interviews also surface conditions under which the mechanisms' downstream consequences are contained. Three protective conditions (institutional slack, preparatory heuristic, conditional approval) are visible in the episode evidence already presented, most

5. Findings

clearly in EP-02 and EP-03. The interviews reveal a fourth, interpersonal rather than procedural, condition.

Respondent A names it directly and quantifies its weight in his own role:

Een informele lijn, dat jij als contractmanager gewoon weet: OK, somebody has my back. ... Ik denk dat mijn werk voor 85% alleen maar op dat relationele zit. En dat trekt het project er doorheen.

[Respondent A, CL-021 and CL-017]

Respondent B describes the same condition from outside the case complex and estimates its scale: most changes are handled informally before they ever reach the formal track.

Wat je ziet is dat heel veel wijzigingen, of eigenlijk 99% van de wijzigingen, worden uitgedeeld in de bouwkeet.

[Respondent B, CL-109]

Three interviews converging on this condition, from inside the case complex, from outside it, and from an adjacent project, support labelling it relational backing and treating it as Strong.

Two interviews document boundary conditions under which relational backing degrades. Respondent C describes a workload boundary:

Vooral aan de voorkant was er zo'n slechte samenwerking tussen project-management en contractmanagement en constructie. En dat kwam omdat mensen overloaded waren met werk. Ik heb met mijn project manager bijna nog nooit meegekeken op wijzigingen van mij.

[Respondent C, CL-216]

Respondent D describes a physical-distance boundary: hybrid working arrangements above a certain threshold degrade the informal coordination the condition depends on (CL-311). Respondent E documents a third failure mode at governance level: senior signatories who are not engaged during the assessment cycle raise questions that should have been resolved earlier, at the point formal sign-off is requested:

Je hebt een soort partijen die op directieniveau allemaal moeten tekenen. Maar die worden niet allemaal goed aan de voorkant meegenomen en geïnformeerd. Dus die gaan eigenlijk pas, als het ter besluitvorming wordt voorgelegd, de kritieke vragen stellen. En dan zeggen ze: 'is dat afgestemd met onze mensen?'. Terwijl je bent één projectteam, hoezo 'onze mensen'?

[Respondent E, CL-412]

Each boundary condition rests on a single interview and is individually Preliminary, though the three converge analytically on the same requirement: unstructured time, physical proximity, and engaged authority among the principal cross-role actors.

Three further candidate protective conditions were named in the interviews, each Preliminary. Respondent B requires the contractor to demonstrate the contractual basis for a change before formal assessment begins, an anti-B1 device operating upstream of the CIA cycle (CL-116). Respondent D describes a change control board as a multi-view bias counterweight, deliberately exposing the decision to actors with different biases (CL-307). Respondent D also describes the worst case position theory, an internal pre-approval strategy, taking the worst plausible cost position before negotiation begins so that neither contractor pressure nor an internal approval deadline can be used as leverage:

What we're doing now is we're just working on a basis that we take the worst position to get the financial approval. (...) Then we still negotiate. So we're not under pressure with the contractor or with the internal pressure of the financial approval.

[Respondent D, CL-318]

5.4.5. Institutional patterns beyond the per-change mechanisms

Two further patterns recur across the interviews but operate at a different level than the three per-change mechanisms: one concerns the counterpart's behavior rather than the assessor's, and one concerns the institutional setting within which both operate.

The first is strategic delay. Four respondents, working under four different contractual regimes, an integrated-contract standard with no fixed deadline (INT-02), a standard with a 14-day notice period (INT-03), FIDIC with a 28-day notice period (INT-04), and a price-offer-expiry mechanism (INT-05), independently describe the same pattern: the contractor delays its response on the formal track until the employer's internal approval deadline forces a settlement on terms the contractor prefers. Unlike B1 to B3, which describe bounded rationality induced in the assessing actor, this pattern describes bounded rationality induced in that actor by a counterparty exploiting the same time-pressure mechanism from the opposite side. After four converging interviews, this inter-organizational form of bounded rationality is treated as Strong.

The second is a procedural-fiction observation. Four respondents describe the formally specified closure procedure as rarely followed as written in practice; the contract standards leave the closure moment of the commercial track institutionally undefined, creating the space within which both the per-change mechanisms ([subsection 5.4.1](#)–[subsection 5.4.3](#)) and strategic delay operate. After four independent observations under four contractual regimes, this is treated as Strong and reported as a precondition for the per-change mechanisms rather than as a separate mechanism in its own right.

Table 5.3.: Bounded rationality strategies by episode

	B1	B2	B3	Primary evidence summary	B1 weight	B2 weight	B3 weight
EP-01	Yes	Yes	Yes	Combined early warning with first approval; coarse cost heuristic; supplier negotiation still open at closure	Strong (n=5)	Prelim. (n=1)	Prelim. (n=1)
EP-02	No	No	No	Phased, deliberative cadence; reserve-cable heuristic embedded as proactive control. Contrast case.	Prelim. (n=1)	Prelim. (n=1)	Evidence of absence
EP-03	Partial	Partial	Yes	Second approval taken under unknown partner pricing; effectuation conditioned on later price and terms	Not coded	Not coded	Prelim. (n=1)
EP-04	Yes	Yes	Yes	First approval taken with consequences still to be quantified; schedule-recovery heuristic at second approval	Strong (n=3)	Not coded	Prelim. (n=1)
EP-05	Yes	Yes	Yes	Two-by-two scenario matrix as substitute for full assessment; combined approval under cable deadline	Not coded	Prelim. (n=1)	Moderate (n=2)
EP-06	Yes	Yes	Yes	Discipline silos miss use-phase impact; board asked to commit before contractor pricing exists	Prelim. (n=1)	Not coded	Prelim. (n=1)

Together, these three strategies produce the CIA completeness outcomes examined in the next section: observable limitations in the scope and depth of assessment that, in several episodes, allowed certain impacts to remain latent at the moment of decision closure.

5.5. CIA completeness outcomes: scope and depth

This section answers the first part of SQ4 by characterizing the scope and depth of CIA in each episode. Drawing on the framework developed in [section 3.5](#), each episode is positioned within the scope and depth matrix ([Figure 3.1](#)), and the evidence for that positioning is described. The section focuses on what was and what was not assessed; the explanation of why, in terms of the strategies documented in [section 5.4](#), is developed in [chapter 6](#).

[Figure 5.1](#) plots the six CIA episodes on the scope and depth matrix introduced in [section 3.5](#). Episodes are colour-coded by the status of downstream corrective work. Only EP-02 sits in the deep-and-broad quadrant, supporting its role as the analytical contrast case. EP-04 reaches the broadest formal consultation in the sample but does not translate that breadth into analytical depth, illustrating the broad-but-shallow pattern. EP-01 and EP-06 occupy the broad-but-shallow region under different conditions: EP-01 at the moment of the first approval, EP-06 cumulatively as new stakeholders entered the assessment over its 617-day cycle. EP-05 is the lowest-depth episode of the sample, but it is not a narrow one: eight roles contributed, and every impact category was touched, placing EP-05 at the far broad end of the scope axis. Compression of the approval moments therefore did not narrow what was looked at, only how far the reasoning was carried. Taken together, four of the six episodes fall in the broad-but-shallow quadrant, while the narrow-and-shallow quadrant, the configuration the matrix identifies as highest-risk, is empty in this sample. The observed failure mode is not that too little was consulted, but that what was consulted was not reasoned through. EP-03 sits within a narrow scope but at moderate depth: its closure was contingent on later partner pricing and terms, thereby embedding the conditional approval protective condition. The two contrast cases (EP-02, EP-03) are the only episodes without downstream rework. They support the framework claim developed in [section 5.6](#): downstream rework is not a direct measure of CIA quality but of the configuration of constraints, closure mechanisms, and protective conditions operating around the CIA cycle.

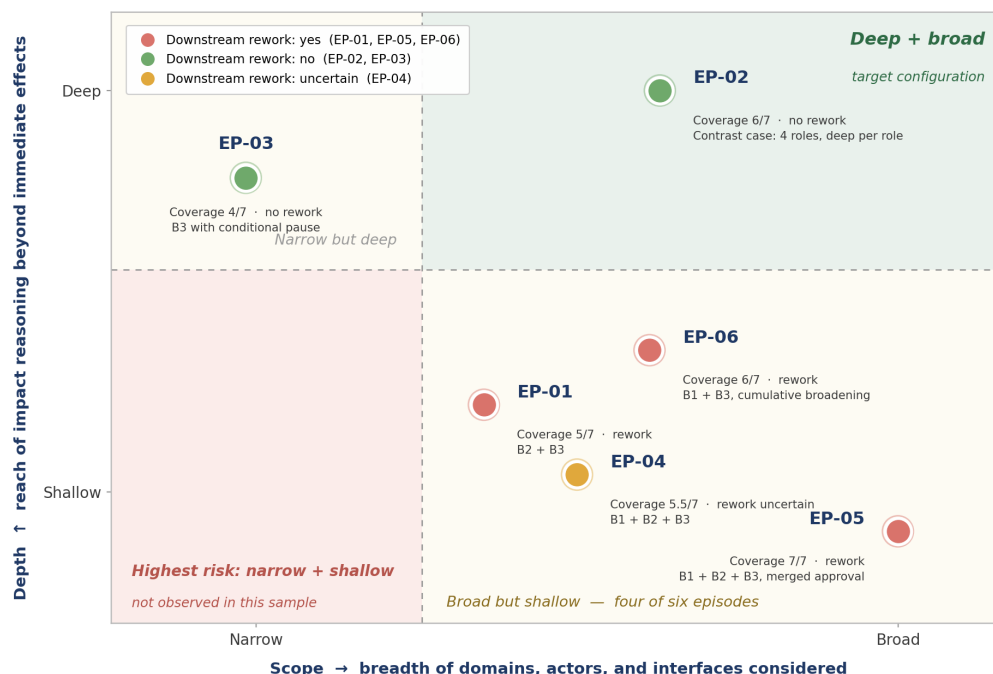


Figure 5.1.: CIA episodes positioned on the scope and depth matrix

5.5.1. Scope: breadth of domains, actors, and interfaces considered

Scope specifies which domains, actors, and interfaces were formally included in the assessment. The six episodes vary considerably on this dimension. EP-04 shows the broadest formal scope, with nine disciplinary roles plus external regulators consulted, although, as noted in [section 5.4](#), the breadth was not matched by analytical depth. EP-02 shows the opposite configuration. Four roles contributed, the narrowest disciplinary range in the sample, yet each contributed substantively, and the assessment reached six of the seven impact categories, a coverage exceeded only by EP-05. Breadth in EP-02 was therefore obtained per contributor rather than per participant: the iterative refinement structure under Contract B left time to interrogate the contributions that were already present instead of adding new ones. This is what places EP-02 in the deep quadrant of [Figure 5.1](#), and it is the reason its scope position should be read as adequate rather than wide. EP-03 shows the narrowest formal scope, with five roles and the lowest category coverage in the sample, reflecting a compressed cadence and a concentration of decision-making authority within the steering group. EP-05 is the counter-example to the assumption that compression narrows scope: despite a merged approval under acute external time pressure, eight roles contributed and every impact category was touched. EP-01 occupies a middle position. EP-06 is analytically distinctive: the formal scope at any single decision point was modest, but the cumulative scope across 617 days became broad as new stakeholders, notably landowners, entered the assessment retrospectively, after the initial CIA had closed. Scope, therefore, is not a single-moment property of the assessment but a cumulative one that depends on how open the assessment remains to later stakeholder entry.

5.5.2. Depth: reach of impact reasoning

Depth characterizes how far impact reasoning extends beyond immediate effects. Three depth-failure patterns appear across the sample.

The first is the downstream-rework blind spot. EP-01, EP-05, and EP-06 show assessments that addressed the immediate construction-phase impacts but did not trace impacts into later phases: the cost of rework, mobilization overhead, and use-phase agricultural impacts, respectively.

The second is the interface-economics blind spot. EP-03 and EP-05 addressed the technical interface but did not fully trace the contractual or financial implications of the closure.

The third is the lifecycle blind spot, seen most clearly in EP-06. The construction-phase CIA addressed the transport and execution requirements, but did not test these against the operational requirements that emerged once the asset was handed over. EP-02 is again the analytical contrast: the assessment explicitly tested the change against the lifecycle requirements, so as to prevent the connection having to be taken out of service later, supporting placement in the deep-and-broad quadrant. Depth is related to the contractual context: deep assessments are concentrated in Contract B; Contract A produces variable depth depending on the time-pressure profile; and the Contract C and D interfaces produce systematically thinner depth on lifecycle considerations.

The documented scope and depth limitations establish the completeness profiles for each CIA episode. The next section examines which of these limitations later became visible as downstream corrective work.

5.6. Downstream manifestations: corrective work

This section addresses the second part of SQ4 by documenting the corrective work that occurred after the assessed changes were implemented and tracing plausible connections between CIA scope and depth limitations and specific downstream outcomes. I treat corrective work here as a potential indicator of latent impacts, consequences that did not surface during CIA, rather than as a direct measure of CIA quality or of individual error. [Table 5.4](#) gives an overview of downstream outcomes by episode.

5.6.1. Episodes with corrective work

Four episodes (EP-01, EP-04, EP-05, EP-06) show downstream corrective work in some form. EP-01 produced a redesign (D1) consisting of duplicate design effort once the external permitting authority's

Table 5.4.: Downstream corrective work by episode

	Corrective work?	Type	CIA limitation linked to	Plausible link
EP-01	Yes	D1 Redesign	Narrow stakeholder scope at first approval; permitting authority excluded	Time-pressure-driven early closure and coarse cost heuristic produced duplicate design effort once the regulator's terms materialized.
EP-02	No	D6 None	None; analytical contrast case	Low operational time pressure permitted a full interface analysis; the reserve-cable heuristic embedded a proactive control. Demonstrates the contra-pole: low constraints, high completeness, no rework.
EP-03	No	D6 None	Closure conditioned on later information	Early closure on the internal track was made operationally safe by an explicit waiting condition for partner pricing and terms: conditional approval as a protective condition.
EP-04	Uncertain	D6 None	Sunk-cost encapsulation and stranded design path	Early closure produced a five-month access delay, absorbed at the second approval by buying schedule recovery (€99k) rather than physical rework.
EP-05	Yes	D1 Redesign	Schedule protection over interface completeness	Cable production proceeded before the municipal ruling; the eventual negative ruling forced a return to original locations, with €190k mobilization overhead and partial civil-works rework.
EP-06	Yes	D2 Resequencing	Lifecycle blind spot; use phase ignored during construction-phase CIA	Discipline silos translated transport requirements without testing the operational agricultural interface; downstream effort to lower 14 road locations and add a temporary transport route (approximately €218k).

terms materialized after the first approval. The linked CIA limitation was the formal exclusion of the permitting authority's interface from the initial scope, traceable to the combination of operational time pressure (a four-week external deadline) and a coarse cost heuristic. EP-04 produced a subtler form of corrective work: a five-month access delay was absorbed at the second approval through a €99k schedule-recovery heuristic, rather than through physical rework. The financial premium is itself a downstream cost of the early closure. EP-05 produced the most consequential corrective work in the sample: the negative municipal ruling forced a return to the original cable locations, requiring approximately €190k in mobilization overhead and partial rework of completed civil works. EP-06 produced the most extensive corrective work in terms of time: 14 road locations had to be lowered after the agricultural impact during the use phase became apparent, and an additional temporary transport route was required during the rework period. The documented cost is approximately €218k, and the social and relational cost exceeded that amount.

5.6.2. Contrastive episode: no corrective work

EP-02 is the analytical contrast case for the chapter. The combination of low operational time pressure, manageable information incompleteness, and a phased decision architecture under Contract B enabled an integrated interface analysis that considered the component's full lifecycle. The downstream outcome is the complementary contrast to EP-05 and EP-06: no physical rework, no schedule-recovery premium, and the proactive insertion of a reserve element that pre-empts a class of operational risk that would otherwise have materialized over the 25-year service life. EP-03 is a secondary contrast case: despite early closure (B3) at the second approval, the documented condition that effectuation would wait until the partner's price and terms were known embedded a conditional approval protective condition that prevented the closure from resulting in downstream rework. Together, the two contrasting cases support the framework claim that downstream corrective work is not a direct measure of CIA quality but rather a measure of the configuration of constraints, closure mechanisms, and protective conditions operating around the CIA cycle.

5.6.3. Cross-episode pattern

Across the six episodes, the relationship between CIA scope and depth and downstream corrective work is mediated by three factors. First, the contractual environment matters: Contract B produced the only deep-and-broad assessment in the sample and the only no-rework outcome. Second, the closure mechanism matters: combined approval closures systematically coincided with downstream rework or schedule-recovery premiums. Third, the protective-condition configuration matters: when conditional approval was embedded in the closure (EP-03), the closure did not result in rework, even though the formal CIA depth was modest. These three patterns are developed in [section 6.1](#) as the empirical basis for the layered reading of construction-phase bounded rationality.

6

Discussion

This chapter discusses the findings from [chapter 5](#) within the bounded rationality framework outlined in [chapter 3](#), refined through six semi-structured interviews. [section 6.1](#) interprets which of the three mechanisms operate under which constraints, introduces a bounded-rationality layer preceding the CIA cycle, and adds a fourth protective condition not visible in the documents. [section 6.2](#) discusses implications for CIA theory and its relation to the literature reviewed in [chapter 2](#). [section 6.3](#) sets out the leverage points that are developed into tools in [chapter 7](#). [section 6.4](#) details what the interviews confirm, what they suggest, and where the analysis is limited. Evidence weights follow the calibration rules specified in [subsection 4.5.5](#). The six interviews are INT-01 (Respondent A), INT-02 (Respondent B), INT-03 (Respondent C), INT-04 (Respondent D), INT-05 (Respondent E), and INT-06 (Respondent F).

6.1. Interpreting findings through bounded rationality

6.1.1. Time pressure as the dominant trigger

[subsection 5.3.1](#) establishes, at Strong evidence weight, that time pressure functions as the dominant precipitating constraint across the case complex, rather than one of three parallel triggers. The range across the sample makes the point concrete. EP-05 sat under a cascade of external deadlines terminating in a cable production date with no slack, and shows every one of the three mechanisms operating at once. EP-02, assessed over a window of roughly six months, shows none of them. Between those poles the pattern is not one of proportional degradation but of a threshold: once the window closes below what deliberation requires, the mode of reasoning changes in kind rather than in degree.

This finding sharpens the framework presented in [chapter 3](#), which introduced time pressure, information uncertainty, and cognitive load without ranking them, and it is consistent with Kahneman (2003) dual-process account: under compressed decision windows, System 1 processing dominates not merely at the margin but categorically, determining whether an actor enters a bounded rational regime at all. Information uncertainty and cognitive load then operate as conditioning factors that shape which of the three mechanisms specified in [section 3.3](#), truncated search, heuristic substitution, or early closure, becomes key once inside that regime, rather than as independent triggers of comparable weight.

This triggered hierarchy is not anticipated by the CIA literature reviewed in [section 2.3](#). Toolkits such as Sun et al. (2006) fuzzy-logic change framework, and multi-criteria approaches such as Gunduz and Mohammad (2020) RII and AHP method, implicitly treat time, information, and cognitive constraints as parallel inputs to an assessment-quality function: more of any one degrades quality by a comparable increment. The empirical pattern here suggests a different structure a single dominant trigger that switches the actor's mode of reasoning, inside which the other constraints determine form rather than degree. This has a direct implication for CIA tool design, developed in [chapter 7](#): tools intended to operate under time pressure should be evaluated against the regime they will actually be used in, not against an idealized, pressure-free baseline.

Promoting time pressure to a dominant trigger sharpens a static reading of the framework but does not yet capture the dynamics by which the regime sustains itself. Three feedback paths visible in the

6. Discussion

empirical material refine the framework into a process reading. First, time pressure compresses the information-gathering phase and therefore raises uncertainty at the decision moment; EP-01 provides a documentary instance in which an order-of-magnitude estimate produced under time pressure was accepted as a decision basis, leaving the actual cost uncertain at closure (CL-015). Second, the combined burden of compressed time and elevated uncertainty raises cognitive load, because unknown unknowns are cognitively expensive to manage; the mechanism responses that follow are predictable consequences of overloaded analytical capacity. Third, those responses close decisions with residual uncertainty, which consumes downstream time when the residue surfaces as corrective work. The B3b portfolio-settlement finding is the visible endpoint of this third path, at which actors close an entire backlog of partly resolved decisions because the cycle has become unmanageable. The triggered hierarchy is therefore not just a ranking; it is a process that escalates itself in the absence of protective conditions. The three feedback paths and their connection back to time pressure are shown in [Figure 6.1](#), which refines the static conceptual model of [section 3.6](#) into a process model of bounded-rational escalation.




Figure 6.1.: Refined conceptual model with feedback dynamics

6. Discussion

The EP-01 instance discussed above traces one full pass through the loop shown here: a compressed information-gathering phase produced an order-of-magnitude cost estimate that was accepted as the decision basis, leaving the actual cost uncertain at closure.

6.1.2. Bounded rationality is partly inherited from the contracting stage

[subsection 5.3.4](#) establishes that a substantial share of the constraint intensity observed within construction-phase CIA episodes originates upstream of the CIA cycle itself, at the contracting stage, via a design-maturity pathway (Strong) and a procurement-fragmentation pathway (Moderate). The episodes cannot show this directly, and that is itself informative: none of the six originates at the contracting stage, so the condition is invisible from within the episode record even where its consequences are not. EP-06 is the clearest case of those consequences. The use-phase requirements for the road entries had never been set at all, and the assessment could not have retrieved information that did not yet exist. What presents within the episode as a depth failure is, upstream, a design-maturity decision taken before the episode began.

This finding extends the CIA literature's implicit scope in an important way. Existing CIA approaches, whether traceability-oriented (Arnold & Bohner, [1993](#)), BIM- and graph-based (Isaac & Navon, [2009](#), [2013](#)), or uncertainty-aware (Sun et al., [2006](#)), are without exception framed as tools that operate within the construction phase, downstream of design and contracting decisions. None of these approaches, nor the change-causation literature reviewed in [section 2.2](#) Muhsen and Rezouki ([2022](#)) and Keane et al. ([2010](#)), treats the contracting stage itself as a source of the constraints that later limit CIA's scope and depth.

This upstream framing also complicates a common implicit assumption in the rework literature: that corrective work observed downstream reflects a failure of construction-phase assessment (Forcada et al., [2017](#); Love et al., [2022](#)). Where the design-maturity pathway is active, part of what later surfaces as corrective work may be more accurately understood as an already-known, deliberately deferred cost, decided at the contracting stage on the rational ground that the cost of staged delay exceeded the anticipated cost of downstream rework, rather than as evidence that construction-phase CIA underperformed. This is consistent with the distinction Love et al. ([2022](#)) draw between avoidable and unavoidable corrective work, but sharpens it: some corrective work that appears unavoidable from within the construction-phase CIA episode is not unavoidable in principle. It is pre-committed at a stage the CIA literature does not examine.

The label meta-BR captures this structure: a bounded-rationality decision taken at the contracting stage, on design maturity or on procurement structure, that pre-shapes the bounded-rationality environment the construction-phase CIA team subsequently inherits. This layering, a bounded-rationality decision producing the conditions for further such decisions downstream, is not addressed by Simon ([1955](#)) original single-episode account of satisficing, nor by the mechanism triad set out in [section 3.3](#), both of which examine bounded rationality within a single decision rather than across nested decision stages. It is, however, anticipated in outline by the organizational extension reviewed in [subsection 2.4.4](#), in which routines established in advance determine which questions a later assessment will ask. The meta-BR layer is therefore best read as an extension of the bounded-rationality literature to multi-stage decision chains rather than a challenge to it, and it identifies a leverage point, developed in [subsection 6.3.1](#) and as Tool 4 in [subsection 7.3.1](#), that sits outside the reach of any construction-phase CIA tool, however well designed.

6.1.3. Per-change mechanisms: B1 truncated search and B2 heuristic substitution

[subsection 5.4.1](#) and [subsection 5.4.2](#) establish, at Strong and Preliminary evidence weights respectively, why the two per-change mechanisms occur and how each is enacted. EP-04 shows why the relational reading matters. Nine disciplinary roles contributed to the assessment memo, several returning no more than an acknowledgment that they had no comments. Read as a diligence failure, this is nine people not doing their jobs. Read through the account developed below, it is nine people correctly inferring that a substantive answer was not what the consultation was actually asking for, and answering the question they were posed rather than the one the framework would have posed.

Both mechanisms are consistent with Simon (1955) original account of satisficing as an economizing response to search cost, but the interview evidence specifies the currency of that cost more precisely than Simon's framework does on its own. For B1, the cost is not only temporal but relational: Respondent A's account (subsection 5.4.1) reframes truncated search as a trust-preservation strategy rather than merely a time-saving one, a dimension the CIA-support literature (Gunduz & Mohammad, 2020; Sun et al., 2006) does not anticipate, since it treats incomplete verification purely as an information-processing shortfall.

For B2, Kahneman (2003) dual-process account explains why a difficult question gets replaced with an easier one, but it does not predict which easier question will be substituted. The four sub-types identified here, pick-your-battles, internal-narrative, risk-minimization, and confidence-substitution, show that the substitute is not arbitrary: it tracks whichever proxy is cheapest to defend within the actor's own accountability structure, respectively relational standing, negotiation narrative, risk exposure, and source credibility. The internal-narrative sub-type in particular resonates with Muhsen and Rezouki (2022) observation that stakeholders may contest root-cause attribution (section 2.2); here the contestation is anticipated and pre-empted through selective framing before it can occur. Read together, the two mechanisms suggest that CIA support tools which treat B1 and B2 as failures of diligence, rather than as structured and defensible responses to specific accountability pressures, are likely to be resisted rather than adopted, a point developed further in the tool design logic of chapter 7.

6.1.4. Per-change early closure (B3a) and portfolio-level early closure (B3b)

subsection 5.4.3 establishes two refinements to early closure: a second B3a vehicle, site-instruction backfill (Strong), alongside the combined approval, and a second scale of closure entirely, portfolio-level bundled settlement (B3b, Moderate). EP-03 and EP-05 make the stakes of the distinction concrete, and they are the sharpest pair in the sample precisely because so much is held constant between them. Both closed early. Both closed under high time pressure with a material input unresolved, partner pricing in EP-03, and a municipal ruling in EP-05. EP-03 produced no corrective work; EP-05 produced roughly €190k in mobilization overhead and the partial demolition of completed civil works. The mechanism was the same, and the constraint profile was near-identical. What differed was that EP-03's closure named the outstanding item and specified the condition on which it would reopen, and EP-05's did not.

Both refinements extend the CIA literature reviewed in section 2.3 beyond its current scope. Traceability- and dependency-based approaches (Arnold & Bohner, 1993; Isaac & Navon, 2009, 2013) are, without exception, specified at the level of a single change: they model how one modification propagates through one project topology. Neither the site-instruction-backfill finding nor the B3b finding fits that unit of analysis. The former shows that closure can precede assessment altogether, under a safety and substance distinction the existing literature does not draw; the latter shows that closure can also operate across a portfolio of changes simultaneously, a scale the CIA literature has not examined.

The B3b finding also gives an empirical mechanism to a pattern the rework literature describes only in aggregate. Ibbs et al. (1998) note that cumulative effects may occur when multiple changes interact over time, but do not specify a decision mechanism that produces cumulative handling. Portfolio-level closure is that mechanism: it is not a deviation from CIA practice but a rational response to the carrying cost of separate handling once the change population exceeds a manageable size, consistent with Simon (1955) satisficing logic applied one level up, from a single decision to a decision population. This reframes downstream corrective work following a bundled settlement not necessarily as evidence that any individual CIA episode failed, but as a possible consequence of a portfolio-level decision to stop evaluating changes individually, a distinction the case-level rework literature (Forcada et al., 2017; Love et al., 2022) does not currently draw.

6.1.5. Protective conditions: a fourth layer

subsection 5.4.4 establishes a fourth protective condition, relational backing, alongside the three derived in section 3.4, together with three documented boundary conditions under which it degrades. EP-02 shows the conditions operating in combination rather than singly. A phased approval architecture supplied the institutional slack, the reserve-element heuristic functioned as a prepared reference rather than an improvised proxy, and no closure was taken before the service-life analysis had been

completed. It is the only episode in the sample in which all three procedural conditions held simultaneously, and the only one with a deep-and-broad assessment.

This finding addresses a gap in the CIA-support literature reviewed in [section 2.3](#): every approach surveyed there, whether traceability-based (Arnold & Bohner, 1993), BIM- and graph-based (Isaac & Navon, 2009), fuzzy-logic (Sun et al., 2006), or multi-criteria (Gunduz & Mohammad, 2020), is procedural or computational in character. None specifies a relational precondition for its own operability. The relational backing finding suggests this is a meaningful omission rather than a matter of scope: the interview evidence indicates that procedural safeguards function only when a non-procedural, interpersonal condition is intact, and degrade in step with it under workload or reduced physical proximity.

This connects to a point made only briefly in the complexity literature reviewed in [section 2.1](#): that project complexity has a socio-technical character, arising from how knowledge and responsibility are distributed across people, not only from technical architecture (Lafhaj et al., 2024; National Academy of Construction, 2022). Relational backing operationalizes that socio-technical claim specifically for CIA. It identifies the interpersonal capital that allows the social half of socio-technical complexity to be managed fast enough to keep pace with the technical half. Where CIA-support research has asked how to represent technical interdependence more completely, this finding suggests an equally consequential question the literature has not asked: how to sustain the relational infrastructure that makes any representation usable under time pressure at all.

6.2. Implications for CIA theory

The framework that emerges from the documentary analysis and six interviews differs from the framework presented in [chapter 3](#) in five respects, each of which extends the existing CIA literature reviewed in [section 2.3](#).

The first refinement is a triggered hierarchy among bounded-rationality conditions. Time pressure is the dominant precipitating condition, while information uncertainty and cognitive load are conditioning factors whose mechanisms become key inside the regime. Existing CIA tools treat the three as comparable inputs to an assessment-quality function. The empirical pattern suggests this misallocates analytical attention. Time pressure does not merely degrade assessment quality at the margin; it determines whether the assessment enters a bounded rational regime at all. CIA support tools that aim to operate under high time pressure should be evaluated against this baseline, not against an idealized pressure-free reference.

The second refinement is the meta-BR layer that sits upstream of the construction-phase CIA cycle. Two pathways feed into it: a design-maturity pathway, in the form of strategic closure on FEED maturity at contracting (Strong), and a procurement-fragmentation pathway, in which the project is split into many small contracts under FID price pressure (Moderate after corroboration from INT-03 and INT-06). Both produce the same downstream consequence: the construction-phase CIA team inherits a configuration in which the per-change mechanisms become salient. The CIA literature treats CIA as a downstream activity to be supported with better tools. The interviews suggest that the most consequential leverage points are upstream of the CIA cycle, at the contracting stage, where CIA support tools cannot reach during the construction phase.

The third refinement concerns the scale at which early closure operates. The existing literature on CIA limitations addresses per-change closure, in which an individual change assessment is shortened. The interviews surface a second scale, portfolio-level closure via bundled settlement. This is not a deviation from CIA practice. It is the rational management of CIA portfolio overhead when the carrying cost of separate handling exceeds the marginal value of doing so. The distinction between B3a and B3b extends the analytical reach of the framework from individual decisions to decision portfolios, and it explains a pattern of practice that has been largely absent from the CIA literature.

The fourth refinement concerns the protective-condition layer. The CIA-support literature has emphasized procedural and institutional safeguards, checklists, dependency models, and multi-criteria templates, but has underdeveloped the relational dimension. The six interviews show that the operability of procedural safeguards under sustained time pressure depends on a non-procedural condition: pre-established interpersonal capital between counterparts. This is the relational backing protective condition, and it is the meta-condition under which the other three remain operable. CIA support cannot be reduced to tool design; it must include attention to the relational infrastructure within which tools

are used.

A fifth refinement extends the constraint set from a static triple of independent inputs to a dynamic system with feedback amplification. The three constraints, time pressure, information uncertainty, and cognitive capacity, amplify each other along three identifiable paths: time pressure compresses the information-gathering phase and raises uncertainty; compressed time plus elevated uncertainty raise cognitive load; and the mechanism responses close decisions with residual uncertainty that consumes downstream time. The CIA literature largely treats the three constraints as parallel inputs to an assessment-quality function. The empirical pattern across the six episodes and six interviews suggests instead that they function as a self-reinforcing process: once the regime is entered, each decision reduces the resources available for the next, until either a protective condition breaks the loop or the actor closes at the portfolio scale. This dynamic refinement is what makes the B3a and B3b distinction analytically intelligible. Per-change closure is the local consequence of the feedback loop reaching the decision moment; portfolio-level closure is the systemic consequence of the loop having compounded across many decisions. Both are predictable outcomes of the dynamic; neither is a deviation from CIA practice but a rational response to its accumulated load. The dynamic refinement is captured visually in [Figure 6.1](#).

Two further findings, strategic delay and the procedural-fiction observation ([subsection 5.4.5](#)), sit outside the per-change mechanism scheme and warrant separate theoretical placement. The contract-management literature reviewed in [section 2.2](#) documents disputes and coordination breakdowns as by-products of ambiguous or contested change causation (Keane et al., 2010; Muhsen & Rezouki, 2022) but treats them as friction, not as a resource one party can deliberately exploit. Strategic delay reframes this: where a counterpart's incentive structure rewards waiting, bounded rationality is not only a constraint the assessor suffers under but a mechanism a counterpart can deliberately induce. This has no analogue in the CIA literature reviewed in [section 2.3](#), which treats the assessing party as the sole locus of bounded rationality, but it is precisely what the contractual extension reviewed in [subsection 2.4.4](#) predicts: where one party cannot foresee every contingency and the other knows it, the resulting incompleteness becomes a strategic resource rather than a technical gap.

The procedural-fiction observation, in turn, complicates the adaptability-gap critique already raised in [subsection 2.3.5](#) (Sun et al., 2006; Woods & Dekker, 2000). That critique frames the gap between formal procedure and dynamic site reality as a tooling problem, in which models fail to keep pace with a changing site. The interview evidence suggests the gap is, at least in part, structural rather than technical: the contract standards themselves leave the closure moment undefined, so no tool operating within the formal procedure can fully close it. This qualifies the practical contribution developed in [chapter 7](#): procedural tools can improve what happens inside the formal track, but cannot, on their own, resolve a gap the contractual framework leaves open by design.

Taken together, these contributions present construction-phase bounded rationality as a layered phenomenon rather than as a flat set of three mechanisms. The foundational layer is meta-BR at the contracting stage. The next layer is institutional: the procedural-fiction observation that contract standards leave the closure moment undefined. The middle layer is the three per-change mechanisms, now refined with the B3a and B3b distinction, an expanded B2 typology of two documentary and four interview sub-types, and the safety and substance distinction within B3a. The surface layer is the four-condition protective profile, with relational backing as the meta-condition that sustains the other three. Each refinement can be tied to the episode that most clearly demonstrates it: the triggered hierarchy to the contrast between EP-05 and EP-02, the meta-BR layer to EP-06's unset use-phase requirements, the scale distinction to EP-03 against EP-05, the protective-condition layer to EP-02, and the dynamic reading to the EP-01 sequence traced in [subsection 6.1.1](#). The refinements are not five separate observations but five aspects of the same six decision histories, read at different levels of resolution. This layered structure is the principal theoretical contribution of the thesis.

6.3. Implications for practice

The layered reading of bounded rationality identifies four leverage points at which CIA support can intervene. Each operates at a different layer, and each implies a different type of intervention. The concrete tool designs that follow from these leverage points are developed in [chapter 7](#); this section sets out the leverage logic only.

6.3.1. Upstream: front-end design maturity (against meta-BR)

Where strategic closure at the contracting stage produces a low-maturity design package, downstream CIA support cannot recover the loss. The most consequential intervention is upstream: organizational discipline at FEED gate review, supported by a rework-risk budget that is visible to the decision authority that authorizes the start of construction. This leverage point lies upstream of the construction-phase CIA cycle, which limits what construction-phase tools can do about it. The practical recommendation is therefore that CIA support tools report on residual design-maturity risk at change initiation, so that the share of CIA effort attributable to inherited incompleteness becomes visible rather than absorbed silently. This leverage point is operationalized as Tool 4 in [subsection 7.3.1](#).

6.3.2. Per-change: pre-formal alignment and conditional approval (against B3a)

Per-change early closure is most effectively mitigated by a structured pre-formal alignment step before the first approval. The step rehearses disagreement informally so that the formal approval meeting either confirms a closure or defers it with explicit conditions, rather than producing closure under pressure. The intervention is procedural but depends on the relational condition: it cannot be mandated as a formal step without losing the informal character that makes it effective. Conditional approval, visible in EP-03, operates on the same leverage point from a different angle: closure on the internal track is taken while the decision is explicitly conditioned on later information.

Respondent A described pre-estimation as an anti-closure procedure in INT-01: producing one's own cost estimate before the contractor's offer arrives, so that the closure decision can be defended on substance. Respondent E describes the institutionalized form of the same condition, an independent cost-expertise reviewer backed by a market database:

In het huidige project hebben we een kostendeskundige die meekijkt. Die heeft heel veel ervaring in wat iets kost en heel veel gegevens over zijn projecten verzameld, dus die heeft een up-to-date database en die kan zeggen 'dit is marktconform' en kan ook een raming maken tegenover de kostenopstelling van de aannemer. Dat is een power tool om de aannemersopstelling te challengen, niet alleen maar te kunnen zeggen 'ik vind het te duur', maar op inhoud challengen.

[Respondent E, CL-410]

The independent cost-reviewer role institutionalizes the preparatory-heuristic condition and turns it from an ad-hoc per-change exercise into a staffable function. Two analytical implications follow. First, the operational form makes the condition reproducible across organizations and projects rather than dependent on the personal initiative of a senior actor. Second, the reviewer is also an anti-bundling device: independent reference prices do not depend on the contractor's own cost trace, so the bundling tactic loses its leverage. The intervention is operationalized as Tools 1, 2, and 3 in [section 7.2](#).

6.3.3. Portfolio: periodic review (around B3b)

Where portfolio-level closure via settlement is anticipated, the appropriate support is a periodic portfolio review that surfaces which classes of change have been accumulating and what aggregate risk they represent. The aim is not to prevent the settlement, since bundling is often the rational management choice, but to ensure that it is taken with an informed view of the portfolio risk profile rather than as a uniform clean slate. Below the tolerance threshold described by Respondent B, roughly 10% of contract value, no review is realistic; above it, the review becomes the institutional venue at which B3b is taken deliberately rather than by default. This leverage point is operationalized as Tool 5 in [subsection 7.3.2](#).

6.3.4. Cross-cutting: relational investment (sustaining the protective layer)

Relational backing operates as a meta-condition for the other three. Its operability depends on prior investment in interpersonal capital across the employer-contractor interface and within each organization. This is the hardest leverage point to operationalize because it cannot be purchased; it must

be cultivated over project lifetimes. The interviews surface two boundary conditions: the condition degrades under workload (INT-03) and under hybrid working arrangements that reduce physical proximity (INT-04). The practical implication is modest but consequential. Roles whose function depends on relational backing, notably the contract manager and the lead change facilitator, should be staffed with continuity in mind, and project transitions that disrupt the relational layer should be recognized as CIA risk events in their own right. This leverage point is operationalized as Tool 6 in [subsection 7.3.3](#).

The four leverage points follow with different evidence weights. The first two are well supported by the documentary evidence and converging interviews. The third rests on the B3b finding (Moderate, INT-01 and INT-02). The fourth rests on the relational backing condition (Strong, three interviews). The concrete tool designs in [chapter 7](#) are developed at the same evidence weights, and each tool is presented with explicit conditions under which it is most and least effective.

6.4. Limitations

Four limitations qualify the findings of this thesis. The first concerns the case base. The empirical analysis rests on a single case complex, with six episodes drawn from a longlist of 32 produced by the case-selection procedure described in [section 4.3](#). The within-case sample is sufficient to identify cross-episode patterns, but generalization to other case complexes, particularly those with different contractual environments or different design-maturity profiles at start of construction, remains an empirical question. Triangulation across two case-complex roles, two cross-project procedural experts, and two adjacent-case practitioners addresses this in part, but does not substitute for evidence from other case complexes.

A second limitation that complements the case-base limitation concerns the documentation environment of the case complex. The case organization operates a structured information-management system that captures project documentation in a consistently traceable form. This is a feature of the case complex, not of construction-phase CIA in general. In projects without comparable documentation systems, the audit trail used in this thesis would be substantially harder to assemble, and the visibility of the mechanisms would correspondingly degrade. The implication for the findings is directional rather than corrective: the mechanisms identified in [chapter 5](#) and refined in this chapter are visible in a most-favorable documentation environment. In environments with weaker documentation, the same mechanisms are presumably present but less visible, and their downstream consequences presumably more severe rather than less. The findings should therefore be read as a floor rather than as a typical case: construction-phase bounded rationality in this case complex is conservative evidence for construction-phase bounded rationality in general.

The third limitation concerns the documentary record. Document analysis can only capture what was written down. The combination of formal and informal practice that the interviews describe, the informal track, the strategic-delay behavior, and the worst-case position strategy, is, by construction, underrepresented in the documentary record. The interview-based findings on the informal track therefore depend on the interview evidence for corroboration rather than on the documents already analyzed. This is a structural feature of studying informal practice through any single-source method, not a flaw to be remedied within the case complex.

The fourth limitation concerns the weight of the new claims. Most of the framework refinements developed in [section 6.1](#) and [section 6.2](#) rest on two or more interviews together with the documentary evidence, at Strong or Moderate weight. A smaller set of candidate findings rests on a single interview each and is reported as Preliminary throughout: the four interview-sourced B2 sub-types (pick-your-battles, internal-narrative, risk-minimization, and confidence-substitution), the required contractual grounding condition, the worst-case position strategy, and the change control board condition. Each is internally consistent with the documentary evidence but is not corroborated by it. The thesis treats these as candidate findings that further interviews are designed to test. If further interviews disprove them, the framework returns to the layered reading without those specific elements, while retaining the Strong and Moderate findings that rest on multi-source evidence.

7

Practical Framework and Recommendations

This chapter translates the empirical findings and theoretical interpretation developed in [chapter 5](#) and [chapter 6](#) into a practical support perspective for construction-phase CIA. [section 7.1](#) presents the overarching framework logic and the design principles that govern the tools. [section 7.2](#) provides three point-of-use tools, one for each per-change bounded rationality mechanism, and [section 7.3](#) provides three cross-cycle interventions operating above the level of a single decision. [section 7.4](#) describes how the six are carried into an operational environment.

7.1. A bounded-rational CIA support perspective


This section sets out the perspective from which the practical contributions in [section 7.2](#) and [section 7.3](#) were developed. The perspective follows directly from the layered reading of construction-phase bounded rationality developed in [section 6.1](#). CIA support is viewed as a series of interventions functioning within a CIA cycle. The cycle's constraints, time pressure first, information uncertainty second, and cognitive load third, are not to be eliminated but managed.

The design principles consist of three key points. The first is constraint-realism: support tools are evaluated for what they enable under high time pressure rather than under an idealized pressure-free reference. The second is leverage asymmetry: interventions earlier in the layered structure, such as front-end design maturity and contracting choices, yield larger downstream benefits than interventions later in the structure, such as per-change procedure design, but they are also harder to deploy within an ongoing project. The third is relational integrity: support tools that disrupt the informal track on which procedural CIA in practice depends can be net negative even when their formal logic is sound.

The six interventions that follow operate at two different levels, which the discussion below keeps distinct. Tools 1 to 3 ([section 7.2](#)) are point-of-use interventions: tools a Change Manager or steering group applies within a single CIA cycle. Tools 4 to 6 ([section 7.3](#)) are cross-cycle interventions: staffing and governance policies that operate above the level of any single decision, targeting the meta-BR layer, the portfolio scale, and the relational infrastructure respectively. [Figure 7.1](#) places the six interventions on the CIA cycle and shows the distinction the two levels rest on: the three point-of-use tools fire at a specific moment within a single change, at the same time the three standing conditions are set up once and require no action per change. Both levels follow from the layered framework in [section 6.1](#), but only the first has the form of a tool in the tools sense; the second is better understood as the organizational precondition under which the first can keep working.

Inside one change cycle — three tools

used by the Change Manager, reviewed by the Contract Manager



Owned by the Programme Director. No action required per change.

Figure 7.1.: When each intervention fires within the CIA cycle

7.2. Tools inside a single change cycle

This section presents the three point-of-use tools, one for each per-change bounded rationality mechanism. Each is specified at the level of trigger, ownership and decision rule; the operational thresholds referred to below are contract-specific and require calibration in deployment. The moment at which each tool fires is shown in [Figure 7.1](#).

7.2.1. Tool 1: Depth Check (While the assessment is being written)

The depth check is a single-page form embedded in the change initiation workflow, completed by the Change Manager and reviewed by the Contract Manager. It asks three questions: which disciplines have been consulted; what the depth of each consultation was, classified as a substantive ask, a sanity check, or an information-share; and, for each consultation logged as either of the latter two, what would have to change for it to become a substantive ask. It activates when the consultation phase closes, before the first approval, and again before the second approval if the depth profile has not changed.

The operational core is the consequence attached to the answer. Where the profile is broad but shallow across more than an organizationally set proportion of consultations, the closure cannot take the combined-approval route, which forces the depth question into the choice of closure format rather than leaving it as a matter of judgment. The failure mode this design anticipates is the form being completed as paperwork; the third question is what prevents it, because a perfunctory answer carries a substantive consequence for how the change may then be closed.

The obvious recommendation in this situation is to consult more deeply. The empirical finding that makes the probe worth building is different: the constraint is not willingness but the location of the cognitive load. In EP-01 and EP-04 the consulted roles returned procedural acknowledgment because formulating a substantive ask is work, and that work sits with the Change Manager, not with the consulted discipline. The probe relocates it accordingly. This tool leaves untouched the efficiency benefit of role-based consultation, which under time pressure is often the fastest defensible way to distribute an assessment, and targets only the case where breadth is not matched by content. It is most effective where the consulted role set is stable enough for the Change Manager to build up tacit knowledge of which roles add depth, and least effective where that role set turns over frequently.

7.2.2. Tool 2: Question check at the assessment moment (When the consultation closes)

The question check consists of two metacognitive questions appended to the standard assessment memo template. The first asks what question the assessment is actually answering. The second asks whether that is the framework question or a defensible heuristic, and where it is a heuristic, requires the subtype to be named from the typology established in [subsection 5.4.2](#), standard-pattern, schedule-driving, pick-your-battles, internal-narrative, risk-minimization, or confidence-substitution, together with the reason it is defensible in this instance. The Change Manager drafts, the Contract Manager reads it as part of approval preparation, and the memo cannot be submitted with the second question unanswered.

The decision rule is deliberately permissive. Where a heuristic is named, the memo proceeds; where the framework question is claimed, the assessment depth must support the claim; where the field is left blank, the memo returns to its author. A substitution is flagged for acknowledgment, not for rejection. The failure mode is the reflexive answer that the framework question was addressed, and the requirement to name a sub-type explicitly is what makes that answer costly to give when the assessment does not support it.

The obvious recommendation is to use the framework properly. The finding that makes the audit worth building is that practitioners are already using heuristics deliberately and can defend them, as every sub-type in [subsection 5.4.2](#) was named and defended by the respondent who used it. Treating the framework question as the norm therefore misallocates attention. Consistent with [subsection 5.4.2](#), pick-your-battles and risk-minimization substitutions are typically defensible, and the audit is designed to surface and support them; standard-pattern substitution applied to a genuinely non-comparable prior case, as in EP-04, is the failure the audit exists to catch. The audit requires a reader other than the

memo's author before approval, since a meta-cognitive question answered by the same person who made the substitution collapses into the frame it was meant to surface.

7.2.3. Tool 3: Open-items approval (At the approval itself)

The open-items approval is a formal approval variant that commits the change on the internal track while explicitly conditioning that commitment on later information. The template names four things: what is not yet available, when it is expected, the conditions under which the decision would be reopened, and the cost ceiling above which the conditional approval is automatically void. It is triggered whenever a team faces a closure deadline with a material input still unresolved, and it is positioned in the workflow as a peer of the combined-approval route rather than as an exception requiring justification, the inverse of merging the two approvals into a single early closure. The Contract Manager owns the template, and the Change Manager completes it; both signatures are required.

The operational core is the qualifying rule. Closure on the internal track is permitted only where at least one condition can be specified concretely: a date, a price, a contract clause, or a testable system behavior. Conditions phrased aspirationally, such as subject to further investigation, do not qualify, and where no condition can be specified the decision must be either a full closure or a full deferral.

This rule follows directly from the empirical material. EP-03 was the only episode in the sample in which a conditioned closure was taken, the only one in which the conditions were specified as testable propositions, and one of only two that produced no downstream corrective work. The failure mode the rule prevents is the one the phrase invites: conditions that are rhetorical rather than enforceable, which leave the residue exactly as unrecoverable as an unconditioned closure would have. This tool leaves untouched the adaptive function of early closure itself, which is often the only way to keep production moving, and targets only closure that leaves residual uncertainty unnamed. It is least effective where the conditioning information is genuinely unknowable, such as regulatory uncertainty with no scheduled ruling, since it then degrades into formalized wishful thinking.

7.2.4. The three tools applied to EP-01

[Figure 7.2](#) summarizes the result of applying the three tools to all six episodes; this subsection works through EP-01 in full, because its downstream cost is well documented and its CIA limitation can be traced precisely. A watercourse crossing on Contract A had to be rebuilt as a culvert, and the decision was taken within twenty-three days to meet a fixed external notification deadline. The interface with the authority responsible for the watercourse fell outside the assessment boundary, and once that authority's terms became known, the design work had to be redone.

The depth check (Tool 1) would not have caught this. The tool probes how thoroughly a discipline was consulted once it has been included in the assessment, but the permitting authority was never on the list of consulted disciplines in the first place. A tool that measures the depth of a consultation cannot flag a consultation that never happened. This is a structural limitation of the instrument, not a flaw in how it was calibrated.

The question check (Tool 2) would have caught it. EP-01 is coded as B2 (heuristic substitution): the question of who owned a downstream stake in the change was replaced with the more tractable question of which disciplines fell within Contract A's formal scope. Because the question check requires the assessor to name, in writing, what question the assessment is actually answering, this substitution would have been visible to the person making it at the moment the memo was drafted, precisely the visibility that an unexamined proxy otherwise removes [subsection 2.4.2](#).

The open-items approval (Tool 3) would not have prevented the miss, but it would have limited its cost. The watercourse authority's terms were genuinely unknown at the time of the first approval, and the tool does not require them to be known, only that their absence be logged as an open item, with an expected date and a defined consequence. EP-03, where this was done without any tool prompting it, shows the difference it makes: design work would have proceeded on a record that marked the crossing as provisional rather than settled [subsection 5.4.4](#).

Two observations follow. First, of the three tools, only the question check and the open-items approval would actually have engaged with EP-01; the depth check would not have applied, since the gap was never inside a consultation to begin with. Each of the two applicable tools addresses a different mechanism, so there is no redundancy between them. Second, and more important for the argument in

7. Practical Framework and Recommendations

[section 7.1](#), none of the three tools would have changed the assessment's actual completeness at the time of the decision: the permitting interface would still have sat outside the assessment boundary at first approval, regardless of which tools were in use. What changes is not the scope of the original assessment but whether the resulting gap is documented and its consequences bounded, rather than left to surface as undocumented rework five months later. Across the six episodes, the pattern in [Figure 7.2](#) holds: of the eighteen tool-episode combinations, four are not triggered at all because the mechanism is absent, and one records a practice EP-03 had already adopted without prompting. The set is therefore selective rather than universally applicable, which is what the reach paragraph in [subsection 7.3.4](#) states in general terms.

Point-of-use tools tested against the six episodes

	EP-01	EP-02	EP-03	EP-04	EP-05	EP-06
1 Depth check (consultation depth probe)	not triggered party was never listed	not triggered no B1	not triggered narrow, not shallow	would have caught it nine roles, thin input	would have caught it B1 under merged approval	not triggered silos, each answered fully
2 Question check (substitution audit)	would have caught it cost proxy for stakeholder scope	fires, and confirms reserve element — deliberate	not triggered no B2 coded	would have caught it risk discount; €99k milestone fit	would have caught it 2×2 matrix for siting decision	not triggered no B2 coded
3 Open-items approval (conditional approval template)	would have caught it permitting terms unknown	not triggered nothing unresolved	already done done already, unprompted	would have caught it consequences not quantified	would have caught it ruling not yet known	would have caught it decision before scope

The three standing conditions (inherited-gap report, bundle review, continuity rule) are not shown: they operate above the single decision and are evidenced in the interviews, not in the episode record.

■ Would have caught something the assessment missed ■ Practice already in place; the tool formalises it
■ Fires, and confirms the shortcut was deliberate ■ Not triggered — the mechanism is not present

Figure 7.2.: The three point-of-use tools tested against the six episodes

7.3. Standing conditions around every cycle

The three tools developed in [section 7.2](#) operate within the per-change CIA cycle and target the per-change mechanisms B1, B2, and B3a. [section 6.3](#) identifies three further leverage points that sit outside the per-change cycle: an upstream leverage point at the meta-BR layer, a portfolio leverage point at B3b, and a cross-cutting leverage point that sustains the protective-condition layer. The three interventions developed below operate at those leverage points, and are specified in the same form as [section 7.2](#).

7.3.1. Tool 4: Inherited-gap report (Once per change, read at Director level)

Inherited-gap report is a single field on the change initiation form recording the share of the present change attributable to the residual design-maturity gap at contracting. It is populated by the Change Manager, validated by the Contract Manager, and aggregated at program level. It activates at every change initiation, and the aggregated metric is reported to the Program Director and reviewed at each Director-level project review.

The tool is not corrective, and its design accepts that. The contracting-stage decision has already been taken and cannot be reversed from within the construction phase; what the field changes is whether its downstream cost is attributed or absorbed silently. The decision rule is therefore reportorial rather than gating: no change is blocked by the field, but the aggregate is placed in front of the authority that will take the next contracting-stage decision. The failure mode is that the attribution becomes a routine percentage applied without thought, which the validation step by the Contract Manager is designed to interrupt.

The obvious recommendation here is to improve front-end design maturity, which every party already agrees with in principle. The finding that makes the reporting field worth building is that the design-maturity compromise documented in [subsection 5.3.4](#) was taken deliberately and with full awareness of its consequences, on the rational ground that staged delay cost more than the anticipated rework. A recommendation to be more mature does not engage that reasoning; a visible running total of what the compromise actually cost does. This intervention does not target a mechanism at the point of use, and it does not imply the contracting-stage decision was wrong, only that its cost should be attributed rather than mistaken for a construction-phase assessment failure.

7.3.2. Tool 5: Bundle review (Once a quarter)

The bundle review is a quarterly review of all pending and unresolved changes within a contract, producing a portfolio risk profile categorized by change class, aggregate cost exposure, and tolerance to bundled settlement. The Contract Manager prepares the profile, the Program Director chairs the review, and the settlement decision remains with the Director-level authority holding the contract. It activates above a tolerance threshold of approximately ten percent of contract value, derived from the heuristic described in [subsection 5.4.3](#); below that threshold no review is realistic.

The review produces three outputs: the changes that can be bundled into a clean-slate settlement, the changes that must be excluded because of their individual risk profile, and an aggregate cost ceiling for the settlement. Negotiation then proceeds on those inputs rather than as a uniform handshake. The failure mode visible in the data is categorization collapsing to a single rating, in which every change is treated as bundled; the design feature that prevents it is the requirement that the exclusion category be populated with at least one named change each quarter, with the reason stated.

The obvious recommendation is not to bundle blindly. The finding that makes the review worth building is that portfolio-level closure is a rational response to assessment overhead rather than a deviation from good practice, which is why this intervention explicitly preserves the mechanism it addresses and targets only its indeliberate form, ensuring that the decision to bundle is taken with the portfolio risk profile visible rather than by default.

7.3.3. Tool 6: Continuity rule (Once per project)

The continuity rule is a staffing rule rather than a tool. Three conditions constitute it. The Contract Manager role is held by a single named person for the duration of the construction phase, with a documented

succession protocol where rotation is unavoidable. The Change Manager and Contract Manager are co-located for at least two days a week, with hybrid participation permitted in one-to-one settings but not in the larger group meetings where relational signal degrades. And neither role is loaded above an organizational capacity ceiling without triggering a documented intervention. The Program Director owns the standard; it is audited at every Director-level project review, and the two role-holders flag breaches.

The decision rule is what distinguishes this from a resourcing matter. Where any of the three conditions degrades materially, the degradation is logged as a CIA risk event in its own right, alongside scope and depth, rather than as a staffing issue to be resolved separately. All three failure modes are documented in [subsection 5.4.4](#): rotation breaking continuity, workload above the ceiling degrading cross-role collaboration, and large hybrid meetings failing to sustain the informal channel.

The obvious recommendation is to invest in relationships, and the obvious response is that the organization already does. The finding that makes the standard worth stating formally is that the three procedural conditions the other tools put into practice degrade at the same rate as relational backing degrades. That is what makes it a meta-condition rather than a fourth safeguard, and what makes investment in the per-change tools without investment in the staffing standard a net negative once the standard slips. This is a staffing standard rather than a point-of-use tool, and it protects the conditions the other five depend on rather than targeting any single mechanism.

7.3.4. Cross-cycle boundary

The three cross-cycle interventions share two characteristics with the per-change tools in [section 7.2](#). First, their operability depends on the relational-integrity principle introduced in [section 7.1](#): where the relational layer between counterparts has degraded, the formal logic of each tool remains intact, but its effective reach decreases sharply. Second, they require the documentation environment that the case complex provides. Without that environment, residual design maturity cannot be attributed reproducibly, portfolio-level risk profiles cannot be categorized consistently, and the failure modes of the relational meta-condition cannot be audited against an objective record. The documentation precondition therefore applies to the cross-cycle interventions as it applies to the per-change tools, and is reflected again in the boundary discussion of [section 6.4](#).

7.4. Carrying the tools into an operational environment

The six interventions specified above share a dependency identified in [section 6.4](#) and [section 8.4](#): they require an environment in which decisions, consultations and conditions are captured in a traceable form. The case organization provides such an environment, and this section describes how the interventions were carried into a workflow specification designed for it. The specification is presented as a design, not as a validated tool; deployment and evaluation are identified as a research direction in [section 8.5](#).

The workflow captures a single CIA episode as a structured record whose entity model mirrors the conceptual model of [section 3.6](#) one-to-one. Construction-phase constraints are recorded as properties of the episode; the disciplines, interfaces and actors considered are recorded as related entities carrying an included-or-excluded flag and an exclusion reason, which together yield the scope dimension; identified impacts are recorded individually with an order field distinguishing first-, second- and third-order effects, which yields the depth dimension; and corrective work is recorded against the originating episode through a post-implementation review task raised thirty days after the decision. The record produced is therefore codable against the scheme in [Appendix D](#) with little additional effort, which allows the same instrument to serve as a practitioner aid and as a research data source.

Three design features follow specifically from the findings rather than from general good practice. First, the actor entity records the basis on which each stakeholder was included as either role-based or knowledge-based. This makes the substitution identified as B2 in [subsection 5.4.2](#) visible at the moment it occurs rather than only in retrospect, because an assessment consulting exclusively based on formal role is visibly doing so on the record. Second, the decision entity carries an explicit stopping-rule field distinguishing an explicit rule from an implicit one and from satisficing, which makes the truncation identified as B1 a recorded property of the decision rather than an inference drawn afterwards. Third,

7. Practical Framework and Recommendations

the decision outcome set includes approve-with-conditions as a first-class option alongside approve, reject, and defer, with the conditions themselves as structured data, which is the conditional approval protective condition of [section 3.4](#) expressed as a system affordance rather than as a recommendation.

A prompt library supplies the targeted prompts anticipated in [subsection 1.5.2](#). Prompts are configurable records rather than hard-coded rules, each carrying the strategy it counters, a severity, and an empirical status; only prompts that have been piloted or validated fire in operational use, so the library can be extended as further evidence accumulates without exposing practitioners to untested material. The suppression rules governing the library follow directly from the empirical findings. Prompts that fire too frequently reproduce in software precisely the failure documented in [subsection 5.4.1](#), in which broad procedural consultation generated acknowledgment rather than substance; a prompt that is dismissed unread is the digital equivalent of a no-comments return. Each prompt therefore carries a cool down and a per-episode ceiling, and every dismissal is logged with the assessor's stated reason. Dismissals are treated as data rather than as non-compliance, which both preserves the relational-integrity principle set out in [section 7.1](#) and generates the evidence needed to tune the library over time.

The specification thus put into practices the reframing developed throughout this thesis. It does not attempt to prevent bounded rational assessment, which [subsection 2.4.2](#) and the empirical material both indicate is neither achievable nor uniformly desirable under construction-phase conditions. It makes the residue of that assessment visible and recoverable: what was excluded and why, on what basis consultation was determined, what rule ended the search, and what remained unresolved at closure. This is the practical expression of managed incompleteness rather than completeness as the goal of construction-phase CIA.

8

Conclusion

This chapter answers the research questions, sets out the thesis's theoretical and practical contributions, and identifies directions for further research. It introduces no new empirical claims and no new theoretical arguments. It draws on what has been established in [chapter 5](#) and [chapter 6](#), with explicit reference to the evidence weights developed there.

8.1. Answer to the main research question

The main research question asked how bounded rationality influences the quality of construction-phase Change Impact Assessment. The answer that emerges from the documentary analysis of six episodes and six semi-structured interviews has three components, and it is an ambivalent answer throughout: the same mechanisms that narrow CIA in some episodes also enable functional, timely decision-making in others. A reading of bounded rationality as only a source of limitation would be inconsistent with two of the six episodes and with several of the framework's own protective conditions.

First, time pressure is the primary trigger. It determines whether an actor enters the bounded rational regime at all; information uncertainty and cognitive load operate as conditioning factors that shape which mechanism becomes key inside the regime. This holds regardless of whether the mechanism that follows narrows or protects CIA completeness in the specific episode: time pressure explains when the regime activates, not whether its effect in a given case is costly or adaptive. This finding rests on documentary evidence and three converging interviews (Strong).

Second, the three mechanisms are not uniformly costly. Truncated search, heuristic substitution, and early closure produced the downstream corrective work observed in EP-01, EP-04, EP-05, and EP-06, but the same three mechanisms produced no corrective work at all in EP-02 and EP-03. EP-02's reserve-element heuristic is itself a standard-pattern substitution that pre-empted a class of operational risk across a 25-year service life rather than narrowing the assessment. EP-03's early closure caused no downstream rework precisely because the mechanism was paired with an explicit reopening condition. Several of the interview-sourced B2 sub-types are, in the respondents' own accounts, deliberate and defensible rather than lapses: pick-your-battles substitution knowingly trades assessment depth for relational capital the project will need later, and the worst-case position strategy uses satisficing deliberately, accepting a conservative estimate early specifically to remove a source of externally induced bounded rationality rather than to save time. The mechanisms are therefore best read as adaptive responses that are sometimes, but not always, costly, depending on whether they are paired with a protective condition.

Third, four protective conditions account for the difference between the two outcomes. Three were derived deductively from the framework: institutional slack, preparatory heuristic, and conditional approval, and one was surfaced inductively across three interviews (Strong): relational backing, the pre-established interpersonal capital that allows decisions to be taken outside the formal track without disproportionate personal risk. Relational backing operates as a meta-condition: it sustains the operability of the other three under sustained time pressure, and without it procedural safeguards degrade regardless of how well they are designed on paper. The protective-condition layer is, in effect, the mechanism by which the same three strategies are steered toward their adaptive rather than their costly manifestation.

8. Conclusion

Taken together, these three components describe construction-phase bounded rationality as a layered phenomenon rather than as a flat set of three mechanisms, and as an ambivalent phenomenon rather than a uniformly negative one. The episode with which this thesis opened makes the point in its plainest form. The culvert was assessed correctly and the design was still done twice, not because anyone failed to reason, but because the boundary drawn in twenty-three days did not extend to a party whose terms had not yet been written. Nothing in the three mechanisms would have prevented that. What would have changed the outcome is a record stating that the terms were outstanding, and a design decision that waited on it. This is what makes the practical contribution developed in [chapter 7](#) a matter of selective reinforcement rather than uniform suppression: tools should preserve the adaptive use of these mechanisms while correcting their costly manifestations, a distinction developed for each tool in [section 7.2](#) and [section 7.3](#).

8.2. Answers to sub-questions

8.2.1. SQ1: how is CIA enacted in practice?

Across the six episodes, a common four-stage procedural structure was observed: initiation, assessment, formal authorization, and implementation. The steering group functions as the decision authority across all cases, and the core team analysis memo serves as the central synthesis artifact prepared in advance of each steering group decision. The most analytically significant variation across episodes is the temporal distance between the first and second approval. This distance is partly stratified by contractual context: the phased pattern shows up most clearly under Contract B, while the compressed combined-approval pattern is concentrated in Contract A and at cross-contract interfaces. Two further variations are visible across episodes: the breadth of disciplinary consultation, ranging from four to nine roles, and the total process duration, ranging from three months to over two years, both of which interact with the compressed-versus-phased pattern in ways developed in [section 5.2](#).

8.2.2. SQ2: what constraints characterize CIA episodes?

Time pressure is rated High or Medium in four of six episodes and is the dominant constraint in all four. Information incompleteness is High in five of six. Uncertainty is High in three of six (EP-03, EP-04, EP-05), where interface dynamics were still unfolding at the moment of decision. Constraints co-occur rather than operate in isolation: where uncertainty is present, it tends to amplify time pressure rather than drive the regime independently. Information incompleteness takes three forms across the sample: incomplete cost information, incomplete interface information, and incomplete regulatory information. These forms align with different points in the CIA cycle at which the assessment must close.

8.2.3. SQ3: which strategies are used, and under what conditions?

B1, truncated search, is observed in four of six episodes, with Strong evidence in EP-01 and EP-04 as the broad-but-shallow consultation pattern. It activates under time pressure combined with broad formal consultation requirements: many roles are consulted, but input remains thin. EP-06 shows a different B1 sub-type, discipline silos that miss the cross-domain interface. B2, heuristic substitution, appears in four of six episodes. The documentary evidence surfaces standard-pattern and schedule-driving sub-types; the interviews surface four further Preliminary sub-types: pick-your-battles, internal-narrative, risk-minimization, and confidence-substitution. B3, early closure, appears in five of six episodes, with the combined-approval pattern as the dominant institutional vehicle under Contract A and at cross-contract interfaces. The interviews add a second B3a sub-type, site-instruction backfill, and a second scale of operation: B3b, portfolio-level closure via bundled settlement, observed in two interviews. EP-02 is the positive contrast: low operational time pressure under Contract B produced a deep and broad assessment with B3 positively absent. Two further patterns operate outside the per-change scheme and are visible only in the interview material. Strategic delay, corroborated across four contractual regimes (Strong), is bounded rationality induced in a counterpart rather than self-induced: the contractor withholds its response on the formal track until the employer's own approval deadline forces a settlement

on terms it prefers. The procedural-fiction observation places the mechanisms inside an institutional gap, since the contract standards leave the closure moment of the commercial track undefined.

8.2.4. SQ4: how do strategies shape CIA completeness, and how do missed impacts become visible later?

Across the sample, EP-02 alone occupies the deep-and-broad quadrant; EP-01, EP-04, EP-05 and EP-06 show broad-but-shallow profiles; EP-03 is the only narrow one. The quadrant the matrix identifies as highest-risk, narrow and shallow, is empty in this sample. Three depth-failure patterns recur: the downstream-rework blind spot, where impact reasoning stops at the construction-phase boundary; the interface-economics blind spot, where a technical interface is addressed without tracing its contractual consequences; and the lifecycle blind spot, where use-phase requirements are never tested. Corrective work followed in four of six episodes.

The link between limitation and rework is mediated rather than direct. Three factors condition it: the contractual environment, which produced the sample's only deep-and-broad assessment under Contract B; the closure mechanism, with combined approvals coinciding systematically with rework or a schedule-recovery premium; and the protective-condition configuration, with EP-03's conditioned closure producing no rework despite modest completeness. Downstream corrective work is therefore not a measure of CIA quality but of the configuration operating around the CIA cycle.

8.2.5. SQ5: under what conditions do strategies not reduce CIA completeness?

Four protective conditions were identified. Three follow deductively from the framework: institutional slack, preparatory heuristic, and conditional approval. The fourth, relational backing, was surfaced inductively across three interviews and is the pre-established interpersonal capital that allows decisions to be taken outside the formal track without disproportionate personal risk. The episode evidence shows the conditions operating in both available directions. EP-02 combined low operational time pressure with a phased decision architecture and produced the only deep-and-broad assessment in the sample. EP-03 shows the more analytically interesting case: a bounded rational early closure taken under high time pressure and unresolved partner pricing, which produced no downstream corrective work because the closure was explicitly conditioned on the missing information arriving later. Completeness at the moment of decision was modest; the outcome was nonetheless adequate, because what remained unresolved was named and tracked rather than left implicit.

Relational backing operates as a meta-condition rather than as one safeguard among four: it sustains the operability of the other three under sustained time pressure, and the interviews document three modes by which it degrades, workload above an organizational ceiling, physical distance under hybrid working, and late-stage engagement by senior signatories who were absent during assessment. Where the relational layer holds, procedural safeguards function as designed; where it degrades, they degrade with it regardless of how well they are specified. The condition under which bounded rationality does not reduce CIA completeness is therefore not the absence of the mechanisms, which are present in nearly every episode, but the presence of a structure that makes their residue visible and recoverable.

8.3. Theoretical contribution

The thesis contributes a layered account of construction-phase bounded rationality that extends the existing CIA literature in five respects. One finding underpins the others and is stated first as a premise rather than as a contribution: time pressure operates as the precipitating condition that determines whether an assessment enters a bounded rational regime at all, with information uncertainty and cognitive load conditioning which mechanism becomes key inside it. That decision-makers act under time pressure is not itself a novel observation. What follows from taking it as the baseline condition, rather than as a variable to be reduced, is where this study contributes.

Five contributions follow, each developed against the literature in [section 6.2](#). First, the protective-condition layer: alongside the three procedural conditions derived in [section 3.4](#), the interviews surface

8. Conclusion

a non-procedural condition, relational backing, on which the operability of the other three depends. CIA support therefore cannot be reduced to tool design. Second, the meta-BR layer: strategic closure at the contracting stage, through a design-maturity pathway (Strong) and a procurement-fragmentation pathway (Moderate), produces the environment the construction-phase team inherits, placing the most consequential leverage points upstream of where construction-phase tools can reach.

Third, the scale distinction between per-change closure (B3a) and portfolio-level closure through bundled settlement (B3b), which extends the framework from individual decisions to decision portfolios and names a pattern largely absent from the CIA literature. Fourth, the triggered hierarchy: time pressure determines whether an assessment enters the bounded rational regime at all, while information uncertainty and cognitive load determine which mechanism becomes key inside it. Fifth, the dynamic reading: the three constraints operate not as parallel inputs but as a self-reinforcing system in which each decision reduces the resources available for the next, until a protective condition breaks the loop or the actor closes at portfolio scale. This is what makes the distinction between B3a and B3b intelligible: B3a is the local consequence of that loop reaching a decision moment, B3b the systemic consequence of it compounding across many.

Two findings sit outside the per-change scheme. Strategic delay, observed across four contractual regimes, is bounded rationality induced in a counterpart rather than self-induced, and establishes an inter-organizational category the framework did not previously contain. The procedural-fiction observation locates the mechanisms inside an institutional gap: contract standards leave the closure moment of the commercial track undefined, which is the precondition both for the per-change mechanisms and for strategic delay. The framework's mechanisms are therefore not purely cognitive.

8.4. Practical contribution

The practical contribution operates at three levels: design principles, leverage points, and six specified tools, the last carried into an operational workflow described in [section 7.4](#) but not yet validated in deployment.

Three design principles govern any support intervention for construction-phase CIA. Constraint-realism: tools are evaluated for what they enable under high time pressure rather than under an idealized pressure-free reference. Leverage-asymmetry: interventions earlier in the layered structure, such as front-end design maturity and contracting choices, produce larger downstream benefits than interventions later, but are also harder to deploy within an ongoing project. Relational integrity: tools that disrupt the informal track on which procedural CIA depends can be net-negative even when their formal logic is sound.

Four leverage points follow from the layered framework structure. Upstream, front-end design maturity operates against the meta-BR layer; the most consequential intervention is organizational discipline at FEED gate review. Per-change, pre-formal alignment together with conditional approval operates against B3a; the intervention rehearses disagreement informally so that the formal approval either explicitly confirms or defers, rather than closing under pressure. At portfolio level, periodic review operates around B3b; the aim is not to prevent the settlement but to ensure that it is taken with an informed view of the portfolio risk profile rather than as a uniform clean slate. Cross-cutting, relational investment sustains the protective-condition layer; roles whose function depends on relational backing should be staffed with continuity in mind, and project transitions that disrupt the relational layer should be recognized as CIA risk events.

The six tools split into two levels, consistent with [section 7.1](#). Tools 1 to 3 are point-of-use interventions, applied by a Change Manager or steering group within a single CIA cycle, and each is designed to preserve the adaptive use of its target mechanism while correcting only its costly manifestation: Tool 1 leaves role-based consultation itself untouched and targets only substanceless input; Tool 2 is designed to surface and support defensible substitutions rather than suppress them, flagging only pattern-recognition failures; Tool 3 leaves the adaptive function of early closure untouched and targets only closure that leaves residual uncertainty unnamed. Tools 4 to 6 are cross-cycle interventions operating above the level of any single decision: Tool 4, residual design-maturity reporting, targets the meta-BR layer by making an inherited and already-rational contracting-stage compromise visible rather than absorbed silently; Tool 5, periodic portfolio review, preserves the core rationale of B3b and targets only undeliberate bundling; and Tool 6, the relational continuity standard, is a staffing standard rather

than a point-of-use tool, protecting the conditions on which relational backing depends. All six presuppose that relational backing remains intact, as specified in [subsection 6.1.5](#), and the documentation environment described in [section 7.4](#). In projects without a comparable system, the recommendation is to invest in that infrastructure before deploying the tools it makes operable.

8.5. Recommendations for future research

Four directions follow from the limitations set out in [section 6.4](#).

First, corroboration of the findings reported as Preliminary: the interview-sourced B2 sub-types and the three candidate protective conditions. Each rests on a single source, and a small number of targeted interviews would move several to Moderate. This is the shortest path to strengthening the framework.

Second, comparative testing across contractual environments beyond the four observed here. Two-phase and alliance contracts are the most informative comparators, because the meta-BR layer should operate differently where design responsibility is shared before construction begins, testing whether the design-maturity pathway weakens or whether a different upstream condition replaces it.

Third, a longitudinal study of relational backing. The interviews establish the condition and three of its failure modes, but how it forms, degrades, and is rebuilt across personnel transitions remains open. Tracking it across a full project life-cycle, with attention to handover moments, would show whether the meta-condition can be deliberately cultivated rather than passively inherited.

Fourth, replication in projects without comparable documentation infrastructure. The visibility of the mechanisms in [chapter 5](#) depends on an unusually well-organized documentation environment. A complementary study would test whether the same mechanisms appear when they are harder to trace, and whether weaker documentation widens the interval between closure and the surfacing of latent impacts. The framework predicts both.

Bibliography

- Al Amri, T., & Marey-Pérez, M. (2020). Towards a sustainable construction industry: Delays and cost overrun causes in construction projects of Oman. *Journal of Project Management*, 87–102. <https://doi.org/10.5267/j.jpm.2020.1.001>
- Arnold, R. S., & Bohner, S. A. (1993). Impact analysis: Towards a framework for comparison. *1993 Conference on Software Maintenance*, 292–301. <https://doi.org/10.1109/ICSM.1993.366933>
- Bano, M., Imtiaz, S., Ikram, N., Niazi, M., & Usman, M. (2012). Causes of requirement change: A systematic literature review. *16th International Conference on Evaluation and Assessment in Software Engineering (EASE 2012)*, 22–31. <https://doi.org/10.1049/ic.2012.0003>
- Chen, L., & Whyte, J. (2022). Understanding design change propagation in complex engineering systems using a digital twin and design structure matrix. *Engineering, Construction and Architectural Management*, 29(8), 2950–2975. <https://doi.org/10.1108/ECAM-08-2020-0615>
- Clarkson, P. J., Simons, C., & Eckert, C. (2004). Predicting change propagation in complex design. *Journal of Mechanical Design*, 126(5), 788–797. <https://doi.org/10.1115/1.1765117>
- Construction Industry Institute. (2015). *2015 annual report* (tech. rep.). Construction Industry Institute.
- Forcada, N., Gangolells, M., Casals, M., & Macarulla, M. (2017). Factors affecting rework costs in construction. *Journal of Construction Engineering and Management*, 143(8). [https://doi.org/10.1061/\(ASCE\)CO.1943-7862.0001324](https://doi.org/10.1061/(ASCE)CO.1943-7862.0001324)
- Gigerenzer, G., & Brighton, H. (2009). Homo heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science*, 1(1). <https://doi.org/10.1111/j.1756-8765.2008.01006.x>
- Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology*, 62. <https://doi.org/10.1146/annurev-psych-120709-145346>
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103(4), 650–669.
- Gigerenzer, G., & Selten, R. (Eds.). (2001). *Bounded rationality: The adaptive toolbox*. MIT Press.
- Gräßler, I., Oleff, C., & Preuß, D. (2022). Proactive management of requirement changes in the development of complex technical systems. *Applied Sciences*, 12(4), 1874. <https://doi.org/10.3390/app12041874>
- Gunduz, M., & Mohammad, K. O. (2020). Assessment of change order impact factors on construction project performance using analytic hierarchy process (ahp). *Technological and Economic Development of Economy*, 26(1), 71–85. <https://doi.org/10.3846/tede.2019.11262>
- Hwang, B.-G., Thomas, S. R., Haas, C. T., & Caldas, C. H. (2009). Measuring the impact of rework on construction cost performance. *Journal of Construction Engineering and Management*, 135(3), 187–198. [https://doi.org/10.1061/\(ASCE\)0733-9364\(2009\)135:3\(187\)](https://doi.org/10.1061/(ASCE)0733-9364(2009)135:3(187))
- Ibbs, C. W., Lee, S. A., & Li, M. I. (1998). Fast-tracking's impact on project change. *Project Management Journal*, 29(4), 35–41. <https://doi.org/10.1177/875697289802900405>
- Isaac, S., & Navon, R. (2009). Modeling building projects as a basis for change control. *Automation in Construction*, 18(5), 656–664. <https://doi.org/10.1016/j.autcon.2009.01.001>
- Isaac, S., & Navon, R. (2013). A graph-based model for the identification of the impact of design changes. *Automation in Construction*, 31, 31–40. <https://doi.org/10.1016/j.autcon.2012.11.043>
- Jarratt, T. A. W., Eckert, C. M., Caldwell, N. H. M., & Clarkson, P. J. (2011). Engineering change: An overview and perspective on the literature. *Research in Engineering Design*, 22(2), 103–124. <https://doi.org/10.1007/s00163-010-0097-y>
- Jayatilleke, S., & Lai, R. (2018). A systematic review of requirements change management. *Information and Software Technology*, 93, 163–185. <https://doi.org/10.1016/j.infsof.2017.09.004>
- Kahneman, D. (2003). Maps of bounded rationality: Psychology for behavioral economics. *American Economic Review*, 93(5), 1449–1475. <https://doi.org/10.1257/000282803322655392>
- Kahneman, D., & Frederick, S. (2002). Representativeness revisited: Attribute substitution in intuitive judgment. In T. Gilovich, D. Griffin & D. Kahneman (Eds.), *Heuristics and biases: The psycho-*

- logy of intuitive judgment (pp. 49–81). Cambridge University Press. <https://doi.org/10.1017/CBO9780511808098.004>
- Keane, P., Sertyesilisik, B., & Ross, A. D. (2010). Variations and change orders on construction projects. *Journal of Legal Affairs and Dispute Resolution in Engineering and Construction*, 2(2), 89–96. [https://doi.org/10.1061/\(ASCE\)LA.1943-4170.0000016](https://doi.org/10.1061/(ASCE)LA.1943-4170.0000016)
- Lafhaj, Z., Rebai, S., Albalkhy, W., Hamdi, O., Mossman, A., & Alves Da Costa, A. (2024). Complexity in construction projects: A literature review. *Buildings*, 14(3), 680. <https://doi.org/10.3390/buildings14030680>
- Langley, A. (1999). Strategies for theorizing from process data. *Academy of Management Review*, 24(4), 691–710. <https://doi.org/10.2307/259349>
- Love, P. E. D., Mandal, P., & Li, H. (1999). Determining the causal structure of rework influences in construction. *Construction Management and Economics*, 17(4), 505–517. <https://doi.org/10.1080/014461999371420>
- Love, P. E. D., Matthews, J., & Ika, L. A. (2023). Fast-and-frugal heuristics: An exploration into building an adaptive toolbox to assess the uncertainty of rework. *Production Planning & Control*, 36(3), 251–266. <https://doi.org/10.1080/09537287.2023.2257178>
- Love, P. E. D., Matthews, J., Sing, M. C. P., Porter, S. R., & Fang, W. (2022). State of science: Why does rework occur in construction? what are its consequences? and what can be done to mitigate its occurrence? *Engineering*, 18, 246–258. <https://doi.org/10.1016/j.eng.2022.05.010>
- March, J. G., & Simon, H. A. (1958). *Organizations*. Wiley.
- Muhsen, H. K., & Rezouki, S. E. (2022). The causes influencing the occurrence of variation orders in the construction of buildings. *Journal of Engineering*, 28(8), 34–53. <https://doi.org/10.31026/j.eng.2022.08.03>
- National Academy of Construction. (2022). *Defining project complexity and its sources: Key points* (tech. rep.). National Academy of Construction.
- Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69(1), 99–118.
- Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63(2), 129–138.
- Simon, H. A. (1990). Invariants of human behavior. *Annual Review of Psychology*, 41, 1–19.
- Sosa, M. E., Eppinger, S. D., & Rowles, C. M. (2004). The misalignment of product architecture and organizational structure in complex product development. *Management Science*, 50(12), 1674–1689. <https://doi.org/10.1287/mnsc.1040.0289>
- Sullivan, B. P., Arias Nava, E., Rossi, M., & Terzi, S. (2023). A systematic literature review of changeability in engineering systems along the life cycle. *Journal of Engineering Design*, 34(12), 1046–1098. <https://doi.org/10.1080/09544828.2023.2273248>
- Sun, M., Fleming, A., Senaratne, S., Motawa, I., & Yeoh, M. L. (2006). A change management toolkit for construction projects. *Architectural Engineering and Design Management*, 2(4), 261–271. <https://doi.org/10.1080/17452007.2006.9684621>
- Tracy, S. J. (2010). Qualitative quality: Eight “big-tent” criteria for excellent qualitative research. *Qualitative Inquiry*, 16(10), 837–851. <https://doi.org/10.1177/1077800410383121>
- Williamson, O. E. (1975). *Markets and hierarchies: Analysis and antitrust implications*. Free Press.
- Williamson, O. E. (1985). *The economic institutions of capitalism*. Free Press.
- Woods, D. D., & Dekker, S. (2000). Anticipating the effects of technological change: A new era of dynamics for human factors. *Theoretical Issues in Ergonomics Science*, 1(3), 272–282. <https://doi.org/10.1080/14639220110037452>
- Yin, R. K. (2018). *Case study research and applications: Design and methods* (6th ed.). SAGE Publications.



Interview Guide

Semi-gestructureerde interviewguide

Bounded Rationality in Construction-Phase Change Impact Assessment

Youp Oosterlee · TU Delft en de case-organisatie · v0.4, twee-fase structuur (45–60 min, één episode)

Interview #	
Datum en tijd	
Locatie of medium	
Rol respondent	
Type respondent	☐ Type A (case-betrokken) ☐ Type B (procedurele expert)
Projectreferentie	
Toestemming opname	☐ Ja ☐ Nee
Interviewer	Youp Oosterlee

Preambule en geïnformeerde toestemming

Lees voor of verwoord natuurlijk. Haast deze sectie niet.

Bedankt dat je tijd hebt vrijgemaakt voor dit interview. Mijn naam is Youp Oosterlee en ik ben masterstudent aan de TU Delft. Ik doe mijn afstudeeronderzoek in samenwerking met de case-organisatie. Het onderzoek gaat over hoe wijzigingsbeoordelingen worden uitgevoerd tijdens de uitvoeringsfase van bouw- en engineeringprojecten, specifiek over de beperkingen waarmee besluitvormers te maken hebben en de strategieën die ze gebruiken wanneer tijd beperkt is en informatie onvolledig is.

We voeren het gesprek in twee delen. Eerst praten we wat algemener over hoe wijzigingsbeoordelingen bij jullie in de praktijk verlopen. Daarna zoomen we in op één concrete wijziging uit jouw ervaring.

Er zijn geen goede of foute antwoorden. Ik ben geïnteresseerd in jouw ervaring en jouw perspectief op hoe het in de praktijk werkt, niet hoe het op papier hoort te gaan. Voel je vrij om openhartig te zijn.

Het interview duurt ongeveer 45 tot 60 minuten. Met jouw toestemming wil ik het gesprek opnemen, zodat ik me kan richten op luisteren. De opname wordt getranscribeerd en geanonimiseerd; je naam en projectspecifieke details die jou of de organisatie zouden kunnen identificeren, worden verwijderd voordat materiaal in de thesis wordt gebruikt of gedeeld.

Deelname is volledig vrijwillig. Je mag elke vraag weigeren te beantwoorden en je mag op elk moment stoppen zonder gevolgen.

Heb je vragen voordat we beginnen?

Verkrijg mondelinge of schriftelijke toestemming. Noteer de toestemmingsstatus in de tabel hierboven.

Fase 1 – Algemene procesbeschrijving

Ongeveer 15–20 minuten. Doel: rapport opbouwen en het feitelijke proces in algemene termen in kaart brengen. De respondent praat over wijzigingsbeoordelingen in het algemeen, niet over een

specifieke case. Dat maakt het makkelijker om vrijuit te spreken voordat we inzoomen op één episode in Fase 2.

Q1. Algemene procesbeschrijving

Kun je beschrijven hoe wijzigingsbeoordelingen in de praktijk meestal verlopen tijdens de uitvoeringsfase?

- Welke stappen worden normaal genomen, formeel of informeel?
- Welke documenten of artefacten spelen een rol: wijzigingsformulier, informatieverzoek, tekening, kernteamanalyse?
- Waar wijkt het in de praktijk vaak af van hoe het formeel bedoeld is?

Q2. Betrokkenheid (algemeen)

Wie worden doorgaans betrokken bij zo'n beoordeling?

- Hoe wordt bepaald wie wel of niet betrokken wordt: op basis van rollen, of van wie je weet dat relevante kennis heeft?
- Komt het voor dat relevante mensen niet worden betrokken? Wat is daar meestal de reden van?

Q3. Beperkingen (algemeen)

Welke beperkingen spelen meestal een rol bij dit soort beoordelingen?

- Tijdsdruk: waar komt die vandaan?
- Incomplete informatie: welke gegevens ontbreken vaak op het moment van besluit?
- Onzekerheid: wat is vaak het minst voorspelbaar?
- Optioneel doorvragen: wat wordt daardoor vaak niet gedaan wat je idealiter wél zou doen?

Q4. Besluitvorming: wanneer is het 'voldoende'?

Wanneer wordt een beoordeling doorgaans als voldoende gezien om een besluit te nemen?

- Zijn er vaak nog open vragen op dat moment?
- Wat gebeurt daarmee: gedocumenteerd, opzij gelegd, of meegenomen naar later?

Overgang naar Fase 2

Zachte overgang, geen harde framing. Laat de respondent zelf een episode kiezen; dat vermindert sociale wenselijkheid en geeft een case waar de respondent zich oprecht bij betrokken voelde.

Om dit wat concreter te maken: zou je een recente situatie kunnen beschrijven waarin jij zelf betrokken was bij zo'n wijziging?

Fase 2 – Concrete episode

Ongeveer 25–35 minuten. Doel: de mechanismen en de causale keten in één concrete episode reconstrueren. Alle vragen keren terug naar de gekozen episode. Gebruik de term bounded rationality niet; gebruik gewone-taalequivalenten zoals shortcuts, de beoordeling ingekort, op je ervaring vertrouwd, of het gesprek afgesloten terwijl er nog vragen open stonden.

Q5. Episode en context

Welke wijziging heb je in gedachten? Kun je kort beschrijven wat er speelde en waarom er op dat moment een besluit nodig was?

- Wat maakte deze situatie urgent, of juist niet?
- Was de wijziging meteen duidelijk, of ontwikkelde die zich nog tijdens de beoordeling?
- In welke fase van de uitvoering bevond het project zich, en wie bracht de wijziging als eerste naar voren?

Q6. Proces in deze case

Kun je me stap voor stap meenemen in hoe de beoordeling in deze specifieke situatie verliep, vanaf het moment dat de wijziging werd voorgesteld tot het moment dat een besluit werd genomen?

- Welke stappen zijn ingekort of overgeslagen vergeleken met hoe het normaal zou gaan?
- Waar moest je aannames doen omdat je iets niet zeker wist?
- Welke documenten of informatiebronnen gebruikte je? Waren die direct beschikbaar?

Focus: truncated search — waar stopte de zoektocht naar informatie?

Q7. Betrokkenheid in deze case

Wie zijn er in deze situatie betrokken geweest, en hoe is bepaald wie input moest leveren?

- Was betrokkenheid gebaseerd op formele rollen, of op specifieke kennis van het raakvlak?
- Waren er mensen die achteraf gezien relevante kennis hadden maar niet betrokken zijn?

Focus: heuristic substitution — werd 'wie heeft de kennis?' vervangen door 'wie hoort hier formeel bij?'

Q8. Beperkingen en counterfactual

Met welke beperkingen had je in deze situatie te maken, in termen van tijd, informatie of onzekerheid? En waren er momenten waarop je onzeker was over hoe ver de impact van de wijziging zou doorwerken?

- Hoeveel tijd had je daadwerkelijk om de beoordeling uit te voeren?
- Heeft die onzekerheid beïnvloed wat je wel of niet verder hebt onderzocht?
- **Counterfactual:** als je meer tijd of betere informatie had gehad, wat had je dan anders gedaan?

Focus: counterfactual reasoning — onthult de niet-onderzochte impacts en de grenzen aan diepgang.

Q9. Het 'goed genoeg'-moment

Op welk moment voelde de beoordeling goed genoeg om een besluit te nemen?

- Wat was op dat moment nog onbekend?
- Heb je bewust besloten om bepaalde impacts niet verder te onderzoeken?
- Was er druk om af te sluiten? Denk je achteraf dat het besluit op het juiste moment werd genomen, of te vroeg?

Focus: early closure en satisficing — het kernmoment van bounded rationality in de beslissing.

Q10. Scope, diepgang en downstream effecten

Als je terugkijkt: hoe breed en hoe diep was deze beoordeling? En zijn er na het besluit impacts naar voren gekomen die niet eerder waren geïdentificeerd?

- Welke aspecten zijn duidelijk meegenomen, en welke juist niet of slechts oppervlakkig? Waarom die selectie?
- Werden er later in de uitvoering of operationele fase impacts zichtbaar die in de beoordeling niet aan bod kwamen?
- Kun je die downstream impacts terugleiden naar hoe de beoordeling is uitgevoerd?
- Was er een moment in het proces waarop het anders had kunnen lopen?

Focus: beperkingen in scope en diepgang, en de koppeling tussen beoordeling en herstelwerk.

Behandel het downstream-deel zorgvuldig. Vermijd het kaderen van herstelwerk als gevolg van individuele fouten; kader het als een normaal kenmerk van complexe bouwomgevingen. Als de respondent terughoudend is om uitkomsten toe te schrijven aan beoordelingsbeperkingen, dring dan niet aan.

Q11. Reflectie en verbetering

Als je één ding zou mogen veranderen aan hoe wijzigingsbeoordelingen worden uitgevoerd, gegeven de realiteit van tijdsdruk en onzekerheid, wat zou dat zijn?

- Op welk moment in het proces zou dat verschil maken: bij de start, tijdens de zoektocht, of op het beslismoment?
- Is er iets belangrijks dat we niet hebben besproken en dat volgens jou relevant is voor dit onderzoek?

Afsluiting

Heel erg bedankt voor je tijd en je openheid. Het gesprek is erg waardevol geweest. Ter herinnering: alle informatie wordt geanonimiseerd voordat ze in het onderzoek wordt gebruikt. Als je na het interview vragen hebt, of als je wilt nalezen hoe ik jouw verhaal in de thesis heb gebruikt, neem dan gerust contact met me op.

Bevestig contactgegevens en vervolgstappen, inclusief member checking indien van toepassing. Stop de opname.

Post-interview analytisch memo

In te vullen binnen 24 uur na het interview.

A. Interview Guide

Onderdeel	Vast te leggen
Episode behandeld	Type wijziging en korte omschrijving
Beperkingen waargenomen	Tijdsdruk, onzekerheid, onvolledige informatie; algemeen en case-specifiek
Strategieën waargenomen	Truncated search, heuristic substitution, early closure
Beperkingen in scope	Welke domeinen, actoren of raakvlakken werden uitgesloten?
Beperkingen in diepgang	Hoe ver reikte de impactredenering? Wat werd niet verkend?
Downstream effecten	Type herstelwerk en schijnbare verbanden met beoordelingsbeperkingen
Counterfactual inzichten	Wat had de respondent anders gedaan met meer tijd of informatie?
Vergelijking Fase 1 en 2	Komt het algemene proces overeen met wat in de case is gebeurd, of zit er discrepantie?
Follow-up nodig	Op te vragen documenten, verduidelijkingen, eventuele tweede episode
Opkomende ideeën	Patronen, verrassingen, afwijkingen van theoretische verwachtingen



Interview Transcripts

This appendix documents the interview material on which Chapters 5 and 6 rest and explains why the full transcripts are not reproduced here.

Six semi-structured interviews were conducted between 5 and 25 June 2026. The interview set comprises two case-involved practitioners (INT-01, Respondent A; INT-06, Respondent F), two adjacent-case practitioners (INT-03, Respondent C; INT-04, Respondent D), and two procedural-expert respondents reflecting cross-project (INT-02, Respondent B; INT-05, Respondent E). The interviews were conducted in Dutch (INT-01, INT-02, INT-03, INT-05, INT-06) or English (INT-04), audio-recorded with respondent consent, and transcribed verbatim, including filler words and minor speech-to-text artifacts.

Each transcript was anonymized under a two-layer regime described in §4.7. The first layer was applied by the case organization before transfer of documentary material to the researcher: partner organizations and project owners appear as generic placeholders (Partij X, Partij Y), individuals are referenced through role designations (Contract Manager, Process Coordinator, Environment Manager), and project-specific geographic or contractual details have been masked where they could permit reidentification. The second layer was applied by the researcher before publication: the case organization is referred to as "the case organization" rather than by name, respondents are referred to as Respondent A–F rather than by name, and internal change-record numbers are kept as opaque case markers.

The full anonymized transcripts are stored securely in line with the data-management conditions of the ethical approval obtained from the TU Delft Human Research Ethics Committee (HREC). They are not reproduced in this thesis to preserve the confidentiality of the case organization and the respondents, given the informed consent under which the interviews were conducted.

Verbatim quotes drawn from the transcripts are reproduced anonymously in Chapter 5 and Chapter 6, with each quote referenced by a Code Log identifier (CL-XXX). Each CL-identifier is traceable to a specific timestamp in the corresponding transcript through the Combined Coding Log (Appendix C and the accompanying Excel file). The full transcripts are available to examiners on request, through the supervisory team, subject to confirmation that the request is consistent with the HREC-approved data-management plan.



Coding Log

This appendix summarizes the analytical structure of the coding work. The full per-fragment coding log, comprising 127 coded fragments across the six interviews with verbatim anonymized quotes, sub-pattern labels, and evidence-weight ratings, is provided as a separate spreadsheet and is the authoritative record. This appendix reproduces the three tables most likely to be consulted while reading: an evidence-weight summary by framework component, a cross-interview map of which findings surfaced in which interview, and the register of new candidate findings.

C.1 Evidence-weight summary by framework component

Table C.1.: Evidence-weight summary by framework component

Framework component	Sources	Code IDs	Weight
Time pressure as dominant trigger	Documents (Table 5.2), INT-01, INT-02, INT-03, INT-05	CL-002, CL-027, CL-112, CL-206, CL-416	Strong
B1 truncated information search (broad-but-shallow)	Documents (EP-01, EP-04), INT-01, INT-02, INT-03, INT-05	CL-011, CL-012, CL-015, CL-112, CL-211, CL-212, CL-214, CL-406, CL-408, CL-409	Strong
B2 standard-pattern sub-type	Documents (EP-04, EP-05), INT-01	CL-022, CL-029	Strong
B2 schedule-driving sub-type	Documents (EP-04), INT-01, INT-05	CL-007, CL-416	Strong
B2 pick-your-battles sub-type	INT-01	CL-006	Preliminary
B2 internal-narrative sub-type	INT-02	CL-107	Preliminary
B2 risk-minimization sub-type	INT-04	CL-319	Preliminary
B2 confidence-substitution sub-type	INT-06	CL-506	Preliminary
B3a per-change closure, combined approval	Documents (EP-03, EP-05, EP-06), INT-01	CL-013, CL-024	Strong
B3a site-instruction backfill	INT-03, INT-04 (refines), INT-05	CL-203, CL-204, CL-205, CL-304, CL-414, CL-415	Strong
B3b portfolio-level early closure	INT-01, INT-02	CL-024, CL-025, CL-026, CL-105, CL-106, CL-111	Moderate
Meta-BR design-maturity pathway	INT-01, INT-04	CL-003, CL-004, CL-030, CL-303, CL-309, CL-312	Strong
Meta-BR procurement-fragmentation pathway	INT-03, INT-06	CL-220, CL-221, CL-509	Moderate
Institutional slack	Documents, INT-01	CL-009	Strong
Preparatory heuristic	Documents, INT-01, INT-04 (variant), INT-05	CL-010, CL-019, CL-305, CL-306, CL-410, CL-411	Strong
Conditional approval	Documents (EP-03), INT-01	CL-010 (related)	Moderate
Pre-estimation (independent cost review)	INT-01, INT-05	CL-010, CL-410, CL-411	Moderate
Relational backing (fourth condition)	INT-01, INT-02, INT-03 (failure), INT-04 (co-location), INT-05 (governance failure)	CL-017, CL-018, CL-020, CL-021, CL-109, CL-216, CL-311, CL-314, CL-315, CL-412	Strong
Required contractual grounding	INT-02	CL-116, CL-117	Preliminary
Pre-formal alignment	INT-01, INT-02	CL-019, CL-119	Moderate

continued on next page

Table C.1 continued from previous page

Framework component	Sources	Code IDs	Weight
Worst-case position strategy	INT-04	CL-318	Preliminary
Change control board	INT-04	CL-307	Preliminary
Hiring strategy	INT-04	CL-306	Preliminary
Targeted specialist consultation	INT-04, INT-05	CL-305, CL-413	Moderate
Site-instruction governance threshold	INT-04, INT-05	CL-304, CL-415	Moderate
Adversarial framing, rhetorical pathway	INT-02	CL-104, CL-110, CL-114	Preliminary
Adversarial framing, behavioral (strategic delay)	INT-02 (implied), INT-03, INT-04, INT-05	CL-209, CL-317, CL-404	Strong
Adversarial framing, cause-bundling tactic	INT-05	CL-405	Preliminary
Zero-rationality boundary case	INT-02	CL-108	Preliminary
Procedural fiction across contract standards	INT-02, INT-03, INT-04	CL-118, CL-208, CL-316	Strong
Relational backing failure: workload	INT-03	CL-216	Preliminary
Relational backing failure: physical distance	INT-04	CL-311	Preliminary
Relational backing failure: late-stage director engagement	INT-05	CL-412	Preliminary
Cost, time and quality triad	INT-01, INT-03, INT-04, INT-05	CL-007, CL-218, CL-302, CL-416	Strong

C.2 Cross-interview map

Direction codes: F = first surfaced; C = corroborates an earlier interview; R = refines or adds nuance; B = boundary condition or failure mode; X = contradicts; a full stop indicates the finding was not addressed in that interview.

Table C.2.: Cross-interview map

Finding	INT-01	INT-02	INT-03	INT-04	INT-05	INT-06
Time-pressure dominance	F (CL-027)	C (CL-112)	C (CL-206)	C (implicit)	C (CL-416)	.
B1 broad-but-shallow	C (CL-011, 012, 015)	C (CL-112)	C (CL-211, 212, 214)	.	C (CL-406, 408, 409)	C (CL-503)
B2 standard-pattern	C (CL-022, 029)	.	.	.	.	.
B2 pick-your-battles	F (CL-006)	.	.	.	.	.
B2 internal-narrative	.	F (CL-107)	.	.	.	.
B2 risk-minimization	.	.	.	F (CL-319)	.	.
B2 confidence-substitution	.	.	.	.	.	F (CL-506)
B3a combined approval	C (CL-013)	.	.	.	.	.
B3a site-instruction backfill	.	.	F (CL-203, 204, 205)	R (CL-304)	C (CL-414, 415)	.
B3b portfolio settlement	F (CL-024, 025, 026)	C (CL-105, 106)	.	.	.	.
Meta-BR design-maturity	F (CL-003, 004, 030)	.	.	C (CL-303, 309, 312)	.	.
Meta-BR procurement-fragmentation	.	.	F (CL-220, 221)	.	.	C (CL-509)
Relational backing	F (CL-017–021)	C (CL-109)	B (CL-216, workload)	B (CL-311, co-location)	B (CL-412, gov- ernance)	.
Required contractual grounding	.	F (CL-116, 117)	.	.	.	.
Pre-formal alignment	F (CL-019)	C (CL-119)	.	.	.	.
Worst-case position strategy	.	.	.	F (CL-318)	.	.
Change control board	.	.	.	F (CL-307)	.	.
Hiring strategy	.	.	.	F (CL-306)	.	.
Targeted specialist consultation	.	.	.	F (CL-305)	C (CL-413)	.
Pre-estimation (independent cost review)	F (CL-010)	.	.	.	C (CL-410, 411)	.
Site-instruction governance threshold	.	.	.	F (CL-304)	C (CL-415)	.
Adversarial framing, rhetorical	.	F (CL-104, 110, 114)	.	.	.	.
Adversarial framing, behavioral (delay)	.	F (implied)	C (CL-209)	C (CL-317)	C (CL-404)	.
Adversarial framing, cause-bundling	.	.	.	.	F (CL-405)	.
Procedural fiction	.	F (CL-118)	C (CL-208)	C (CL-316)	.	.
Cost, time and quality triad	F (CL-007)	.	C (CL-218)	C (CL-302)	C (CL-416)	.
Zero-rationality boundary	.	F (CL-108)	.	.	.	.
Tolerance threshold on B3b	.	F (CL-111)	.	.	.	.

C.3 New findings register

Table C.3.: New findings register

Finding	First surfaced	Corroborated by	Weight	Status
Time-pressure dominance	INT-01	INT-02, INT-03, INT-05	Strong	Confirmed, subsection 6.1.1
B3b portfolio-level settlement	INT-01	INT-02	Moderate	Candidate, subsection 6.1.4
B3a site-instruction backfill	INT-03	INT-04 (refines), INT-05	Strong	Confirmed, subsection 6.1.4
Meta-BR design-maturity	INT-01	INT-04 (live case)	Strong	Confirmed, subsection 6.1.2
Meta-BR procurement-fragmentation	INT-03	INT-06	Moderate	Confirmed, subsection 6.1.2
Relational backing as fourth condition	INT-01	INT-02, INT-03, INT-04, INT-05	Strong	Confirmed, subsection 6.1.5
Required contractual grounding	INT-02	—	Preliminary	Candidate, section 6.3
Pre-formal alignment	INT-01	INT-02	Moderate	Refinement, section 6.3
Pre-estimation (independent cost review)	INT-01	INT-05	Moderate	Operationalised, subsection 6.3.2
Targeted specialist consultation	INT-04	INT-05	Moderate	Refinement, section 6.3
Site-instruction governance threshold	INT-04	INT-05	Moderate	Refinement, section 6.3
Adversarial framing, rhetorical	INT-02	—	Preliminary	Candidate, not carried forward
Adversarial framing, behavioral (delay)	INT-02 (implied)	INT-03, INT-04, INT-05	Strong	Confirmed, section 6.2
Adversarial framing, cause-bundling	INT-05	—	Preliminary	Candidate, not carried forward
B2 pick-your-battles sub-type	INT-01	—	Preliminary	Refinement, subsection 5.4.2
B2 internal-narrative sub-type	INT-02	—	Preliminary	Refinement, subsection 5.4.2
B2 risk-minimization sub-type	INT-04	—	Preliminary	Refinement, subsection 5.4.2
B2 confidence-substitution sub-type	INT-06	—	Preliminary	Refinement, subsection 5.4.2
Zero-rationality boundary case	INT-02	—	Preliminary	Candidate, not carried forward
Tolerance threshold on B3b	INT-02	—	Preliminary	Refinement, subsection 6.1.4
Relational backing failure: workload	INT-03	—	Preliminary	Boundary, subsection 5.4.4
Relational backing failure: physical distance	INT-04	—	Preliminary	Boundary, subsection 5.4.4
Relational backing failure: late-stage directors	INT-05	—	Preliminary	Boundary, subsection 5.4.4
Worst-case position strategy	INT-04	—	Preliminary	Candidate, subsection 5.4.4
Change control board	INT-04	—	Preliminary	Candidate, subsection 5.4.4
Hiring strategy	INT-04	—	Preliminary	Candidate, not carried forward
Procedural fiction	INT-02	INT-03, INT-04	Strong	Confirmed, section 6.2
Cost, time and quality triad	INT-01	INT-03, INT-04, INT-05	Strong	Background, section 5.4



Coding Scheme

This appendix defines the codes used in the within-case and cross-interview analysis. The codes are organized by family: constraints (A), bounded rationality mechanisms (B), protective conditions (PC), downstream corrective work (D), context and framing (CTX), and new candidate findings (NEW). For each code, the definition explains what was treated as a positive instance during coding. Where a code was refined or extended during the interview-based analysis, the refinement is noted in the definition. The evidence-weight categories in [Table D.6](#) apply the calibration rules specified in [subsection 4.5.5](#).

D.1 Constraint codes

A-codes are applied to passages that describe the conditions under which a CIA cycle took place. They do not by themselves indicate the presence of a mechanism; they indicate the precipitating conditions under which the mechanisms become likely.

Table D.1.: Constraint codes (A)

Code	Label		Definition
A1	Time pressure		Explicit or implicit deadline pressure within the CIA cycle, originating in production schedules, regulatory deadlines, market deadlines, or contractual milestones. Coded at intensity Low, Medium or High.
A2	Information incompleteness		Information that exists in principle but is not yet available at the moment of decision, such as a cost proposal not yet provided, partner pricing unknown, or a permit ruling pending.
A3	Uncertainty		Information that does not yet exist because the underlying situation has not yet stabilized, such as a regulatory outcome, site dynamics, or a social and relational process still unfolding.

D.2 Mechanism codes

Table D.2.: Bounded rationality mechanism codes (B)

Code	Label		Definition and sub-types
B1	Truncated search	information	Stopping the search for impact information before all plausible effects are explored. Primary manifestation: broad-but-shallow consultation, in which many disciplinary roles are formally consulted but their inputs are thin. Sub-pattern: discipline silos, where cross-domain interfaces are not consulted at all.

continued on next page

Table D.2 continued from previous page

Code	Label	Definition and sub-types
B2	Heuristic substitution	Replacing the framework question, what is the full impact, with an easier proxy. Documented sub-types: (a) standard-pattern, proceeding as in a similar prior change without testing comparability; (b) schedule-driving, substituting the impact question with whether the change fits the milestone envelope. Interview-sourced sub-types, each Preliminary: (c) pick-your-battles, closing a low-impact change on relational grounds; (d) internal-narrative, selectively framing the elements won for internal accountability; (e) risk-minimization, substituting the best option with the lowest-risk option; (f) confidence-substitution, accepting information because it is presented persuasively rather than because it has been verified.
B3	Early decision closure	Committing to a decision before the full information set is in. Operates at two scales, coded separately as B3a and B3b below.
B3a	Per-change early closure	Closure of a single CIA cycle before its full information set is in. Two vehicles: (a) combined approval, in which the first and second approval are merged into a single closure moment; (b) site-instruction backfill, in which execution precedes both assessment and formal closure. Within the second, safety-driven backfill is treated as legitimate and substance-driven backfill as the mechanism the framework tracks.
B3b	Portfolio-level early closure	Closure of many CIA cycles at once through a bundled settlement, taken because the carrying overhead of separate handling exceeds the marginal informational gain. Below a tolerance threshold of approximately ten percent of contract value, portfolio-level closure operates as routine non-attention rather than as deliberate settlement.

D.3 Protective conditions

Protective conditions are institutional, procedural, or relational factors that moderate the mechanisms. [section 3.4](#) derives three deductively; the interview analysis surfaces a fourth at Strong weight, together with a further set of candidates.

Table D.3.: Protective conditions (PC)

Code	Label	Definition and status	Weight
PC-IS	Institutional slack	Available organizational capacity above the minimum needed to execute the formal CIA procedure. Derived deductively in section 3.4 .	Strong
PC-PH	Preparatory heuristic	A prepared and validated reference point available before the decision moment, converting an implicit substitution into a deliberate one. Derived deductively in section 3.4 . Variants observed: independent cost review (INT-01, INT-05) and experience-based judgement (INT-04).	Strong
PC-CA	Conditional approval	A decision taken with explicit conditions that reopen it when the outstanding information arrives. Derived deductively in section 3.4 and documented in EP-03.	Moderate
PC-RB	Relational backing	Pre-established interpersonal capital between counterparts that allows decisions to be taken outside the formal track without disproportionate personal risk. Surfaced inductively; operates as a meta-condition sustaining the operability of the other three. Three documented failure modes: workload (INT-03), physical distance (INT-04), and late-stage director engagement (INT-05).	Strong
PC-PFA	Pre-formal alignment	Rehearsing disagreement informally before the formal approval meeting, so that the meeting either confirms a closure or defers it with explicit conditions. Procedural-informal variant of PC-PH.	Moderate

continued on next page

D. Coding Scheme

Table D.3 continued from previous page

Code	Label	Definition and status	Weight	
PC-PE	Pre-estimation	An independent cost-expertise reviewer with market-database backing who can challenge the contractor's estimate on substance. Institutionalizes PC-PH as a staffable role, and also functions as an anti-bundling device (INT-01, INT-05).	Moderate	
PC-TSC	Targeted specialist consultation	Identifying, per change, the specialist roles whose substantive knowledge is required, rather than routing the change past all formal roles by default. Anti-B1 device (INT-04, INT-05).	Moderate	
PC-SGT	Site-instruction governance threshold	A threshold scheme distinguishing safety-driven from substance-driven backfill of B3a, with supervisor-level instruction authority up to a set value and escalation above it (INT-04, INT-05).	Moderate	
PC-RCG	Required grounding	contractual	Requiring the change initiator to demonstrate the contractual basis for the change before assessment begins. Procedural-formal device operating upstream of the CIA cycle (INT-02).	Preliminary
PC-WC	Worst-case strategy	position	Pre-approving the worst plausible cost internally, so that neither contractor pressure nor an internal approval deadline can be used as leverage. Internal-procedural condition and the operational inverse of satisficing (INT-04).	Preliminary
PC-CCB	Change control board		A multi-view governance body that exposes a decision to actors with different biases, reducing the chance that a single substitution becomes the framing of record. Institutional anti-B2 device (INT-04).	Preliminary
PC-HS	Hiring strategy		Upstream investment in experienced staffing as a precondition for per-decision judgement. Upstream variant of PC-PH (INT-04).	Preliminary

D.4 Downstream corrective work codes

Table D.4.: Downstream corrective work codes (D)

Code	Label	Definition
D1	Redesign	Duplicate or repeated design effort triggered by impacts that surfaced after decision closure.
D2	Resequencing	Re-ordering of construction activities triggered by impacts that surfaced after closure.
D3	Demolition and reconstruction	Physical removal of completed work and its re-execution.
D4	Additional coordination	Substantive coordination effort beyond what the original plan envisaged.
D5	Schedule-recovery premium	Financial premium paid to recover schedule without physical rework, such as an acceleration payment.
D6	No corrective work	The episode produced no observable downstream corrective effort. Used for contrast cases.

D.5 Context and new candidate codes

Table D.5.: Context (CTX) and new candidate (NEW) codes

Code	Label	Definition
CTX	Context and framing	Background, framing, or contextual content that informs interpretation but does not by itself constitute a coded instance of a constraint, mechanism, or protective condition.
NEW	New candidate finding	A pattern surfaced for the first time in any of INT-02 to INT-06 that does not fit cleanly into the deductive scheme. Examples: adversarial framing, the zero-rationality boundary case, the procurement-fragmentation pathway, the worst-case position strategy, and confidence-substitution.

D.6 Evidence-weight categories

Table D.6.: Evidence-weight categories

Category	Definition	Examples
Strong	Three or more coded observations across at least two distinct source types, or sustained corroboration across multiple interviews and the documentary evidence.	Time-pressure dominance (documents and three interviews); B1 broad-but-shallow (documents and three interviews); relational backing (four interviews).
Moderate	Two converging sources: either two interviews from different vantage points, or one interview together with documentary evidence.	B3b portfolio settlement (INT-01 and INT-02); site-instruction backfill refinement (INT-03 and INT-04); pre-formal alignment (INT-01 and INT-02).
Preliminary	A single observation from a single source.	Pick-your-battles sub-type (INT-01); worst-case position strategy (INT-04); confidence-substitution (INT-06).
Evidence of absence	Positive evidence that the mechanism or condition did not operate, as distinct from an absence of observation.	B3 is positively absent in EP-02, the contrast case.
Not coded	No observation was recorded either way. Distinguished from evidence of absence, which is a finding.	Used in Table 5.3 where a mechanism is coded present but no supporting weight was recorded.

D.7 Quote selection criteria

Table D.7.: Quote selection criteria

Criterion	Definition
Representativeness	The quote illustrates a recurring observation across multiple sources or cases rather than an isolated remark. Where several quotes support the same pattern, the one presented is the one whose phrasing most directly captures the mechanism.
Clarity	The wording explicitly shows the mechanism, so that the reader can verify the relationship between code and quote without extensive interpretive scaffolding.
Traceability	Each quote references a specific Coding Log entry and, where applicable, a specific Document Log entry, so that the reader can locate the source and verify the coding decision independently.



Overview of Interviews Conducted

This appendix summarizes the six interviews conducted for this study. The interview guide and the informed-consent preamble read to each respondent are reproduced in [Appendix A](#); the full anonymized transcripts are in [Appendix B](#); and the analytical structure of the coding work is summarized in [Appendix C](#). The per-fragment coding log is held as a separate spreadsheet.

Table E.1.: Overview of interviews conducted

ID	Date	Type	Role	Min.	Lang.	Codes	New findings flagged
INT-01 <i>Resp. A</i>	2026-06-05	A	Senior practitioner in the case complex; 26 years of experience	50	NL	31 (22)	B3b portfolio settlement; relational backing as fourth protective condition; meta-BR design-maturity pathway
INT-02 <i>Resp. B</i>	2026-06-10	B	Senior contract manager with a legal background, approximately 10 years; expert in integrated-contract and FIDIC regimes; outside the case complex	42	NL	21 (16)	Corroborates B3b and relational backing; introduces adversarial framing (rhetorical), required contractual grounding, and the procedural-fiction observation
INT-03 <i>Resp. C</i>	2026-06-10	A (adjacent)	Assistant contract manager and process coordinator on an adjacent industrial project; 80–100 contracts; approximately 3,000 changes	30	NL	22 (17)	Site-instruction backfill sub-type; procurement-fragmentation pathway; strategic-delay tactic
INT-04 <i>Resp. D</i>	2026-06-11	A (adjacent)	Senior project manager with a mechanical-engineering background; shipbuilding and oil and gas; case organization, pre-execution stage	50	EN	20 (14)	Live meta-BR observation; strategic-delay corroboration; worst-case position strategy; B2 risk-minimization sub-type; co-location boundary of relational backing
INT-05 <i>Resp. E</i>	2026-06-22	B	Process consultant in a contract-management role; cross-project experience; external consultancy with an information-management software tool	20	NL	16 (11)	Paper-reality articulation of B1; cause-bundling adversarial pathway; late-stage director failure mode of relational backing; independent cost review as operationalisation of the preparatory heuristic; site-instruction governance threshold
INT-06 <i>Resp. F</i>	2026-06-25	A	Senior mechanical-engineering supervisor in the case complex; prior international pipeline-project experience	30	NL	17 (16)	B1 case-complex instance; B2 confidence-substitution sub-type; procurement-fragmentation corroboration, moving that pathway from Preliminary to Moderate; refinement of the hybrid-working boundary of relational backing

E. Overview of Interviews Conducted

Type A respondents were directly involved in one or more CIA episodes; Type A (adjacent) respondents work on comparable projects outside the case complex; Type B respondents are procedural experts with cross-project experience and no involvement in the case complex. The codes column gives the number of coded fragments, with the number selected for use in the thesis in parentheses.



Declaration on the use of artificial intelligence

What

Anthropic Claude (Opus 4.1 and Sonnet 4.5, accessed through the Claude desktop application, March to September 2026) was used throughout the writing and revision of this thesis.

How and to what extent

The empirical work was carried out without AI. I selected the six episodes from the project record, read and coded the documents by hand, and conducted, transcribed, and coded the interviews myself. I assigned the evidence weights against the calibration rules stated in [subsection 4.5.5](#). No AI tool was given access to the interview material or to the project documents.

I used AI in three ways.

First, as a structural reviewer of drafts. I submitted completed chapters and asked for internal inconsistencies, unsupported cross-references, mismatches between chapters, and claims not supported by the evidence cited for them. This produced several substantive corrections, the most consequential being that the derivation of the protective conditions was missing from [chapter 3](#) although five later cross-references assumed it was present. I checked every correction against the sources and the case material before adopting it.

Second, as an editor and drafter. Some passages were drafted or rewritten with AI assistance, always from material and reasoning I supplied, and always revised by me afterwards. Passages produced through such iterations include [section 3.4](#), [subsection 5.5.1](#), [subsection 7.2.3](#), and parts of [chapter 6](#) and [chapter 8](#).

Third, for the production of figures. [Figure 3.2](#), [Figure 5.1](#), [Figure 6.1](#), [Figure 7.1](#), and [Figure 7.2](#) were generated from plotting code written to my specification with AI assistance, using data taken from my own coding notebook.

Reflection

The tool was most useful for verification rather than generation: comparing a claim in one chapter against the evidence presented in another, locating cross-references that no longer resolved, and counting occurrences across a document too long to hold in mind. It was considerably less reliable on the facts of the case. On two occasions it stated confident conclusions about events that proved to be wrong; in one of these it had conflated two distinct uses of the term "combined approval", and the error surfaced only when I checked the underlying documents. That is why the empirical work was deliberately kept outside the scope of AI assistance: the tool reads an argument well on a second pass, but it does not read a case.

I am responsible for this thesis in full, can account for its content and analytical choices independently, and have checked every claim originating in an AI-assisted draft against its original source.

