Hochschule Neu-Ulm
University of Applied Sciences

Master Thesis

In the Master program
Digital Transformation and Global Entrepreneurship

at University of Applied Sciences Neu-Ulm

Architecturally Hostile: AI Commercialisation, Domain Expertise, and Institutional Logic Clashes in Canadian Residential Construction

1st Examiner: Prof. Dr. Daniel Schallmo

Author: Hunter Sturgeon (Enrollment number: 367963)

Date of Registration: 25.11.2025

Date of Submission: 25.03.2026

Abstract:

The Canadian residential construction sector is currently defined by significant systemic challenges in affordability and a decades-long stagnation in productivity. Despite an influx of investment in Generative and Agentic AI, the industry remains a technological laggard, characterised by a fragmented structure of micro-businesses operating on thin margins. This thesis investigates the barriers and enablers that determine whether entrepreneurial AI innovations successfully navigate the commercialisation interface with Canadian residential construction firms.

Adopting a qualitative, exploratory multi-case study design, the research analyses 12 semi-structured interviews with digital innovators and industry incumbents through the lenses of Institutional Logics, Paradox Theory, Technological Frames, and Sociomateriality. The findings show that the sector is Architecturally Hostile to AI adoption. This hostility is sustained by a Structural Void where fragmentation, Operational Capture, The Moving-Target Trap, and Regulatory Non-Replicability create an inhospitable commercialisation interface.

Crucially, the research discovers a Credibility Threshold empirically grounded in the "Grey Hair Effect", a practitioner-generated theory whereby domain expertise is not merely a competitive advantage but a functional prerequisite for trust and cognitive offloading. To address these systemic failures, the study proposes the **Collaborative Commercialisation Framework (CCF)**. This framework shifts the commercialisation encounter from a competition-driven logic to an interaction-driven co-creation model, providing actionable pathways for startups, incumbents, and policymakers to bridge the structural void between innovation development and viable industry adoption.

Keywords: *AI Commercialisation, Construction Productivity, Institutional Logics, Interaction Interface, Digital Transformation, Digital Entrepreneurship, Canadian Housing.*

AI Usage Statement:

Various Artificial Intelligence (AI) tools were used throughout this thesis. The author practiced *Digital Bricolage* during the writing of this paper, actively using any and all digital tools that could improve various stages of the research process. AI was used to improve working efficiency, as an ideation partner, and to increase clarity and consistency. Primarily, AI was used to articulate and organise written drafts into clear ideas, explain theoretical concepts in plain terms, and as a partner in ideation. AI was consistently used to ensure argumentative clarity and organisation of the thesis was coherent and conformed to academic expectations. Additionally, AI tools were used to transcribe, and after scrubbing them of personal data, identify emerging themes across interviews. Finally, an AI grammar tool was used to check and improve the grammar and ensure uniformity across the document. AI tools and applications were used responsibly and complimentary to the work of the author, and did not serve as a substitute for interpretation or original theoretical contributions.

List of Figures

List of Abbreviations

Framework and Concept Abbreivations:

Technical Abbreviations:

Organisational and Policy Abbreviations:

Table of Contents

1 Introduction

The Canadian housing sector is currently characterised by significant systemic challenges regarding affordability and availability. Despite an expansive geographical landscape and access to raw materials, the sector has arrived at a state of underinvestment and innovation stagnation. CMHC reports that the primary driver of the affordability crisis is a lack of construction productivity rather than speculative investment alone. Between 2019 and 2024, this productivity loss resulted in a national spending increase of 6 to 8 billion dollars, contributing to the approximately 20% increase in new home prices in Canada. This trend is exacerbated by the collapse of productivity growth following the COVID-19 pandemic. (Iorwerth, 2025, p. 2)

The challenge is further complicated by the industry's structure; 69.5% of residential construction companies are micro-businesses operating with low profit margins. These entities often lack the capital required to scale or adopt traditional technological solutions such as robotics or modular construction, which have historically failed to achieve significant market penetration. However, the emergence of Generative and Agentic AI presents a new potential for transformation. With federal investment in construction technology (ConTech) set to increase by 700 million and global investments reaching 6.75 billion in 2025, there is a compelling case for exploring how AI-driven innovation can address these productivity ailments. (Innovation, Science and Economic Development Canada, 2024)

This paper investigates what barriers and enablers determine whether entrepreneurial AI innovations successfully navigate the commercialisation interface with Canadian residential construction firms, and what conditions must be met for adoption to begin. The paper contributes to the academic community by bridging the gap between digital entrepreneurship literature and construction management studies, specifically by theorising the meso-level commercialisation interface as the site where innovation diffusion succeeds or fails. Practically, the research aims to provide the Canadian housing sector – both innovators and incumbents – and its policymakers with a deeper understanding of the diffusion lag. By identifying specific barriers and enablers at the interaction interface, the paper provides a basis for targeted interventions to accelerate adoption. The primary output is a conceptual framework designed to support digital transformation within the sector, accompanied by actionable recommendations for startups, industry practitioners, and policymakers.

2 Literature Review

2.1 The Core Debate: Why does Digital Innovation fail in construction?

2.1.1 Historical Lag: Barriers to Productivity

The residential construction industry has a productivity problem. The Canadian sector's productivity has been stagnant, which cannot be explained only by market conditions - it reflects a structural incapacity for innovation that has compounded over decades and now represents the primary barrier to addressing the affordability crisis. This lack of technology diffusion and productivity growth, particularly when compared to the manufacturing sector, has created the consensus that this industry is a “laggard”. This reputation for conservatism is often quantified by the industry's low aggregate expenditures on research and development (R&D) relative to other industries (McCoy et al., 2008, p. 125). In the Canadian context, the construction sector operates well below full capacity, contributing to a persistent productivity gap (Searle and Rankin, 2025, p. 2). Canadian Residential construction productivity has stayed marginally the same, relative to the majority of Canadian industries, which have shown growth. (Statistics Canada, 2024) while there is some debate that this statistical stagnation may partly result from measurement errors, as traditional economic metrics commonly fail to capture the "service-enhanced" nature of construction or the qualitative improvements in the built environment (Searle and Rankin, 2025, p. 2-3; Winch, 2003). Still, it is very likely that the lack of innovation is responsible for the productivity gap because measurement errors do not account for the divergence between construction and other industries that use the same economic metrics. (Searle and Rankin, 2025, p. 2)

The persistence of this productivity lag is frequently attributed to the industry's unique structural characteristics. Unlike the steady-state environment of manufacturing, construction projects are organised as "loosely coupled systems". A term first introduced by Weick in 1976, that describes a system characterised by fragmentation and temporary relationships (Dubois and Gadde, 2002, pp. 621-623). For residential construction, this means the diffusion of innovation is stifled by the transient nature of project teams; because teams are assembled for unique projects and then disbanded, the long-term collaboration and learning cycles required to refine and scale new processes are difficult to sustain (Bossink, 2004, p. 341; Searle and Rankin, 2025, p. 3). For AI adoption, the impact is two-fold: teams implementing

AI often operate in silos that limit knowledge transfer, while AI agents themselves are constrained by fragmented data that prevents system-wide learning. (Ghimire et al., 2024, p. 20; Onetrace, 2026)

Furthermore, risk aversion in Canadian residential construction is not merely cultural conservatism – it is a rational response to an economic environment where a single failed innovation can threaten firm survival. Builders often view innovation as a financial and liability risk rather than a competitive advantage. There is a prevailing belief that homebuyers and clients prefer "tried and true" products, leading firms to avoid non-traditional technologies to minimise the risk of defects, warranty claims, and call-backs (Koebel et al., 2003, pp. xi, 26; McCoy et al., 2008, p. 125). This risk aversion creates a barrier to diffusion where information channels become siloed; while early adopters utilise diverse information sources (such as universities and trade shows), the majority of the industry relies on a narrow set of sources, slowing the spread of new knowledge (Koebel et al., 2003, pp. 30–31). The thin-margin, survival-driven logic that dominates the sector contributes to this risk aversion, where an innovative misstep or adopting the wrong technology can prove dire for construction firms. For AI adoption, this risk aversion interacts with the pace of improvements in AI technology to create a particularly intractable form of non-adoption. AI technology is iterating quickly, and if a culture of risk aversion results in a wait-and-see approach to AI adoption, by the time the technology is “tried-and-true”, the industry risks arriving too late, lacking both the organisational fluency and data infrastructure that early engagement would have built.

2.2 Existing Explanations for Adoption Failure

2.2.1 Structural Explanations: Fragmentation and B2B

The industry is deeply fragmented, both vertically along the supply chain and horizontally among a vast number of small, competing firms (Koebel et al., 2003, p. 13; Koeleman et al., 2019). Recent data from CMHC indicates that the Canadian residential construction sector is composed of 69.5% micro-businesses (1-4 employees), a structure likened to a "hunting-gathering" era that lacks the financial and human resources to modernise (Laberge, 2025). Residential construction projects are equally fragmented, where decision-making is decentralised, and production is handled by multi-party networks rather than integrated entities (McCoy et al., 2008). Furthermore, Research from the FMI Corporation (a consulting and research group) in 2018 found that 96% of construction data is discarded or uncollected, largely because the fragmented value chain lacks the connectivity required to utilise it (Spare et al., 2018, pp. 2). Consequently, project sites rich in unstructured data remain black boxes for decision-makers. This is particularly challenging for digitalisation efforts like adopting agentic AI, which often requires not just large amounts of data, but connected and consistent cross-project data.

The problem is further exacerbated by the industry's lack of consumer demand input, as is seen in B2C industries. Residential construction operates on a B2B logic where the end-user (the homebuyer/renter) has little direct influence on the process, nor do they frequently possess the technical knowledge to apply pressure on technologies that appear opaque. This insulates the industry from the disruptive "consumer pressure" that accelerates digital transformation in other sectors (Samuelson & Stehn, 2023, pp. 391-392). Without end-user pressure to seek transparency, interoperability, or digital capability, the industry's logic remains oriented toward cost and familiarity rather than innovation – a condition that systematically disadvantages AI solutions whose value compounds over time rather than delivering immediate visible returns.

2.2.2 Variable Regulatory frameworks

A structural barrier largely absent from traditional diffusion models is regulatory heterogeneity. Unlike centralised jurisdictions, Canada's construction sector operates within a multilayered governance regime composed of municipal bylaws, provincial building codes, and federal standards, each of which may impose distinct compliance requirements (RBC, 2024). This institutional fragmentation produces what, for residential construction, can be termed *Regulatory Non-Replicability*: a condition in which solutions validated in one jurisdiction cannot be deployed elsewhere without redesign, recertification, or renegotiation. Regulatory Non-Replicability can be seen as a synthesis of the proximal academic concepts of "Institutional voids" (Khanna & Palepu, 2010), "Replication" (Winter & Szulanski, 2001) and "path dependency" (Arthur, 1989). It offers a justification for why the digital transformation of construction has not yet led to the scaling seen in other Canadian industries, such as software or retail.

Systemic innovations such as AI-enabled project platforms or modular construction systems depend on repeatability to generate economies of scale. However, when approval processes, inspection criteria, and certification standards vary across jurisdictions, the marginal cost of scaling rises sharply. It is difficult for machine-learning to be effective when rules are constantly changing, so even effective AI solutions may only be locally effective because AI systems trained on data from one jurisdiction will be non-compliant with another(Ghimire et al., 2024, p. 19). This is a fundamental limitation of the transferability of AI's core value proposition in construction. What appears to be technological resistance may therefore represent rational economic behaviour; firms decline adoption because innovations cannot be deployed consistently across projects or regions.

Recent policy efforts illustrate recognition of this structural constraint. The One Canadian Economy Act (Bill C-5, 2025) seeks to reduce interprovincial regulatory friction by facilitating labour mobility and harmonising approval pathways for nationally significant projects. From an innovation perspective, such initiatives function not merely as economic policy but as institutional infrastructure for diffusion. By lowering regulatory variance, they allow innovations to be jurisdictionally portable, thus lowering the customisation required for innovators and reducing friction for firms to invest in systemic technologies. When the industry is composed mostly of micro-firms operating with survival-logic, ameliorating the risks associated with innovation adoption is integral for diffusion.

Thus, regulatory fragmentation should not be interpreted as an administrative challenge but as a central explanation for construction's innovation lag. It constrains scaling, suppresses experimentation, and reinforces the short-term decision logic already dominant among micro-firms. This variability is especially consequential at the commercialisation interface, where innovators must demonstrate scalability.

2.2.3 "One-Off" Production and Site Variability

The One-off nature of residential construction functions as a systemic barrier to the data-driven logic required for AI integration. By synthesising the work of Koeleman, Papadonikolaki, and others, it becomes clear that the industry's "one-off" nature creates a structural prevention of data accumulation. While manufacturing relies on the standardisation and repetition of assembly lines, residential construction is defined by a high degree of product and demand variability (Koeleman et al., 2019; Papadonikolaki et al., 2020). This heterogeneity ensures that design and delivery approaches are rarely repeated measurably, depriving AI models of the consistent, high-volume longitudinal datasets they need to "learn" and optimise.

This lack of data continuity is driven by what Samuelson and Stehn (2023, p. 388) identify as "project-based discontinuity." Teams are transient and disbanded upon completion, so the industry lacks a centralised repository for process improvements. In a Canadian context, this discontinuity is exacerbated by extreme site variability – such as unique ground conditions and volatile weather – which forces builders to prioritise "tried and true" methods to mitigate liability (Koebel et al., 2003, pp. xi, 26; McCoy et al., 2008). From an analytical standpoint, these site-specific risks act as "noise" that drowns out the "signal" of innovation; when every project has unique local elements, the enterprise-level data remains too fractured and non-transferable between regions to support the economies of scale that drive AI transformation.

Central to this failure of diffusion is the concept of "hidden innovation" (Ozorhon et al., 2010 pp. 566). In the residential sector, innovation is often sub-systemic or improvised. While these micro-innovations solve immediate project needs, they remain "hidden" because the meso-level interaction interface (the space between the site and the firm) lacks the digital infrastructure to capture them. Without a deliberate framework, AI may actually make hidden innovation *less* likely to diffuse. If AI tools are trained only on standardised data, they may fail to recognise or validate the informal innovations that happen at the site level. Conversely, AI presents a potential solution to surface these innovations. If used as a sensing tool within the meso-layer, capturing onsite messy data, AI could theoretically bridge the project-to-enterprise gap, transforming one-off innovations into permanent firm-level assets. (Ghimire et al., 2024, p. 12).

The convergence of these scholars suggests that the primary obstacle to AI in construction is not the technology itself, but the institutionalised diversity of the projects themselves. To move forward, the industry must transition from viewing projects as isolated events that act as a data black-box to viewing them as iterations that can provide clean data to improve the next project.

2.2.4 “Survival" Logic and Economic Explanations

Understanding the pressures put upon the majority of firms in the Canadian residential construction sector is important for understanding the innovation logic that dominates the industry. Firm size dictates the Canadian industries’ innovation logic. Research by Meng & Brown (2018) found that for small firms, innovation is driven primarily by "survival and prosperity" and "cost savings" (pp. 1217-1218). They lack the "slack resources" of excess capital, time, and prospects to experiment. In contrast, large firms are driven by global competition and have the capacity to adopt long-term strategies that may not see immediate gains, but instead will see improvements later (Kussl & Wald, 2025, p. 382; Meng & Brown, 2018, p. 1218). However, these larger enterprises are much less agile and function on significantly longer time-horizons. Meaning, their barrier to innovation adoption is change management and analysis paralysis. (Meng & Brown, 2018, p. 1218)

Thus, innovators face a bifurcated market, where each segment has different enablers and barriers for AI. Small firms cannot afford, or risk adoption, and large firms are indecisive. For innovators attempting to navigate this market, the bifurcation creates a structural paradox – the firms most operationally suited to benefit from AI adoption lack the resources to engage, while the firms with sufficient resources lack the decisional agility to commit. Without consolidation or significant scaling – similar to the model seen in the Dutch or Swedish housing sectors – Canadian firms struggle to accumulate the resources or infrastructure necessary to adopt AI and automated prefabrication (Laberge, 2025).

2.2.5 Financial and Operational Misalignment

In Canada, the conflict is most acute in the financial structuring of projects, particularly regarding off-site and modular construction. These kinds of startups require significant upfront capital for factory production, often demanding deposits of 30% to 50% of the contract value (Kurup et al., 2024, p. 4). This directly contradicts traditional construction financing, where lenders typically release funds only upon verification of on-site progress (Kurup et al., 2024; RBC, 2024). This disconnect can increase the equity requirement for modular projects by as much as 30% compared to traditional site-built projects (Kurup et al., 2024, p. v). The distinct lack of excess capital within the industry means off-site and modular start-ups struggle to find buyers.

Operationally, the rapid iteration cycles of tech vendors (often 12–24 months or less) are fundamentally disconnected from the long timelines of construction projects, which can span three to ten years from inception to completion (Zachariah & Goldsmith, 2022). For AI technologies, these long industry timelines could prove to be an institutionally structured asset because AI models thrive on sustained data collection. This potential is contingent on incumbents having the means and systems in place to track and report data consistently.

Furthermore, Bridgit, a Canadian construction management platform, offers practitioner insight into how misaligned expectations damage the commercialisation interface, noting that startups that oversell their capabilities to secure market share risk eroding the foundational trust that sustained adoption requires (Bridgit, 2023). For adoption to succeed, innovators must pivot toward effectiveness and simplicity – defined by the ease of understanding and using a product – and align their "push" with the builder's need for immediate risk reduction (Rogers, 2003; Walzer et al., 2025, p. 251). A report from the Business Development Bank of Canada found that successful Canadian examples, such as DOZR, illustrate that adoption occurs when platforms are rooted in "real-world experience" and target specific, tangible inefficiencies like workforce allocation and equipment utilisation (BDC, 2022).

Together, these financial and operational misalignments compound so that even technically viable AI solutions encounter structural hostility at the commercialisation interface – not because the industry incumbents are outright rejecting innovation, but because the economic architecture of construction was built for a world where innovation was never expected to arrive from outside.

2.2.6 The Structural Vacancy: Integration Without an Integrator

The industry lacks a central "systems integrator" – a role common in manufacturing – which creates gaps in accountability and coordination when introducing new technologies, such as novel building envelope systems (Koebel et al., 2003, pp. 10-11). As Prencipe and Tell (2001, p. 1374) identify, the project-based and multi-party structure of construction means that the role of systems integrator falls onto the project, rather than any one individual. The structural absence of a permanent integrating actor has, over time, produced adversarial contractual norms and siloed information flows – adaptive responses to a system that was never designed for sustained cross-party collaboration. (Spare et al., 2018; Zachariah & Goldsmith, 2022). Successful transformation requires a stable structure to enforce the new order. The lack of a clear "process owner" or systems integrator means there is no entity with the authority to stabilise these new digital-social relationships across the supply chain (Samuelson & Stehn, 2023, p. 389). Without this stabiliser to continually ensure innovations become practices, the industry reverts to its default analog practices once the temporary project pressure subsides. This absence of an integrating actor suggests that for AI technology to diffuse, there is a need for a framework capable of coordinating technological, organisational, and regulatory dimensions.

This structural vacancy raises a question the existing literature has not addressed: in the absence of an institutional integrator, what actor might perform this coordinating function at the commercialisation interface – and what characteristics would they need to do so effectively?

2.3 The failures of Existing Theories

To understand why AI adoption lags in construction, it is necessary to interrogate the industry's foundational innovation models. While traditional frameworks categorise innovation by magnitude (Slaughter) or diffusion speed (Rogers), they often fail to account for the unique "systemic" nature of digital tools and the structural rigidity of the construction ecosystem.

2.3.1 Critiquing Slaughter: The Systemic Nature of AI vs.The Fragmented Industry

Slaughter's (1998, p. 226) model of construction innovation, alongside subsequent project-based typologies by Gann and Salter (2000, p. 955) and Hobday (2000, p. 874), serves as the core framework for classifying industry advancements. Within these typologies, digital tools like AI are clearly defined as "systemic" innovations that require the integration of multiple independent components throughout diverse project communities (Slaughter, 1998, p. 226). However, the critical limitation of Slaughter's framework is not purely the observation that systemic innovations clash with fragmented environments. Rather, Slaughter fails because the framework provides no analytical variable for the characteristics of the innovator. It cannot explain why the exact same AI product will successfully embed into a project when championed by a domain expert, yet struggle to gain traction without one.

Slaughter's "system" category inherently assumes the existence of a central "systems integrator" capable of managing coordination across complex, multi-firm alliances (Gann & Salter, 2000, p. 958). Canadian residential construction structurally lacks this permanent capacity. By relying on temporary coalitions, the industry forces the one-off project itself to act as the main mechanism for coordination and integration (Hobday, 2000, p. 874). Consequently, Slaughter's systemic category assumes an integration infrastructure that the residential industry simply cannot generate. Because it assumes a non-existent integrator and ignores the innovator entirely, Slaughter's model cannot reach the operational realities at the meso-level interface. Specifically, it offers no mechanism to analyse how the innovator's characteristics directly impact decision speed, risk allocation, and resource optimisation. When an AI startup introduces a systemic tool, it is the innovator's domain knowledge and interaction strategy – not just the technology's design – that determines whether project resources are optimised, how liability and risk are allocated across the supply chain, and whether the decision speed of the project manager (PM) is supported or hindered.

Ultimately, Slaughter cannot account for innovator characteristics or structural vacancies, and Rogers cannot account for a technology whose characteristics are not fixed at the point of diffusion. Both theoretical failures point to the exact same gap: the unresolved dynamics at the meso-level interface.

2.3.2 Critiquing Rogers DOI Framework:

If Slaughter fails to account for the *actor* required for systemic integration, Rogers' Diffusion of Innovation (DOI) fails to account for the *fluidity* of the innovation itself. While Rogers' (2003)DOI theory has long served as the academic baseline for understanding market permeation (McCoy, 2008, p. 41), recent scholarship argues that the model possesses significant limitations when applied to the complex, fragmented nature of construction digital transformation, particularly fluid innovations like AI technology.

A primary critique of DOI is its inherent "pro-innovation bias" – the assumption that adoption is inherently beneficial and that rejection or "laggardness" is a sign of irrational conservatism (Crofton, 2024; Fichman, 2004; Mohammed, 2023). However, in the case of residential construction, the wait-and-see approach Rogers' framework would diagnose as irrational conservatism, is frequently a rational response to genuine uncertainty. Rogers DOI framework assumes the innovation being diffused is stable, treating it as a "static artifact" (Lyytinen & Damsgaard, 2001, p. 5). The current state of AI technology breaks this framework; AI technology is rapidly developing, and evolves faster than firms' evaluation cycles. For DOI theory, which is built on the premise of assessing relative advantage, compatibility and complexity at a fixed point in time, there is no mechanism for the "moving-target-trap". In the high-risk, low-margin reality of construction, "non-adoption" is frequently a rational economic survival strategy rather than a failure of diffusion (Meng & Brown, 2018, p. 1215; Yang et al., 2025, p. 1). Nevertheless, non-adoption can prove structurally catastrophic if the stability being waited for never arrives.

DOI assumes adopters have the bandwidth to evaluate innovations at the point of diffusion, whereas in residential construction, teams are too consumed with operations for significant R&D. (Zachariah & Goldsmith, 2022)Meaningful evaluation may be structurally unavailable regardless of the potential innovations' impact. This structural bandwidth deficit explains why domain expertise functions as a prerequisite condition at the commercialisation interface rather than merely an advantage - incumbents do not have the R&D capacity, so they have to offload this to innovators. Therefore, the reputation and perceived experience of the innovator are critical because incumbents need to trust that the solution is actually solving a real problem.

A fundamental flaw in applying DOI to the Canadian residential construction sector is its assumption of a relatively homogeneous adopter population making independent, rational choices based solely on an innovation's attributes (Fichman, 2004; Lyytinen & Damsgaard, 2001, p. 5). The Canadian residential construction market is, in reality, deeply bifurcated. On one end, the industry relies on micro-firms operating under a strict, day-to-day survival logic; on the other, large firms frequently suffer from analysis paralysis due to their complex organisational hierarchies and risk aversion. Rogers' DOI provides no mechanism to account for structurally different adoption barriers operating simultaneously within the exact same industry.

Furthermore, DOI treats technologies as stable, standalone artifacts. Lyytinen and Damsgaard (2001, p. 6) forcefully dismantle this "discrete package" fallacy, arguing that complex, networked digital tools do not diffuse as fixed, objective packages from an independent innovator to a passive adopter. Instead, digital innovations are "learning intensive" and "socially constructed," constantly mutating depending on the context in which they are applied (Lyytinen & Damsgaard, 2001, p. 7). Rogers treatment of compatibility is insufficient for the specific demands of socially constructed digital tools. This framework fails to address technologies that require continuous reconfiguration and learning. (McCoy, 2008, p. 47)

Because Rogers views only the innovation itself, DOI ignores the broader environment required for the technology to function. As Adner (2006) points out, successful adoption relies not just on the product, but on the "innovation ecosystem" – the collaborative arrangements through which firms combine offerings into a coherent solution. Adner introduces the concept of "interdependence risk," highlighting the profound uncertainties of coordinating with complementary innovators and systems. In construction, the adoption of systemic tools like AI fails not just because of the technology's inherent characteristics, but because the surrounding ecosystem of data infrastructure, regulatory frameworks, and organisational capacity is not ready to support it. Rogers sees only the innovation; Adner sees the ecosystem. Fichman (2004) echoes this in his critique of the IT innovation diffusion literature, noting that Rogers' dominant paradigm is insufficient for complex technologies because it cannot capture the systemic interdependencies and holistic configurations required for modern digital tools to survive. Consequently, adoption in these complex networks is often driven more by institutional pressures and "herd effects" than by independent, rational evaluations of the tool's features (Lyytinen & Damsgaard, 2001, p. 8).

Rogers cannot account for a technology that will not stabilise, an adopter population that is structurally too consumed with immediate project survival to evaluate it, or a market that is bifurcated by firm size into two different adoption problems. Together with Slaughter’s failure to account for innovator characteristics and the structural vacancy of a systems integrator, this means that the current theoretical toolkit cannot explain what happens at the meso-level interface. Yet, it is precisely at this meso-level interface–where the startup's product logic collides with the builder's project logic–that innovation either succeeds or dies in Canadian residential construction.

2.4 The Meso-Level Gap: Where Innovation Breaks and Theory Falls Silent

2.4.1 Synthesis: The System is Architecturally Hostile

The barriers to AI adoption in Canadian residential construction are not independent frictions that compound additively. To answer the primary research question of this thesis, it is crucial to view these structural barriers as mutually reinforcing features of a complete system. This system was architecturally designed – through fragmentation, project-based transience, survival logic, regulatory heterogeneity, and financial misalignment – for a world where innovation was never expected to arrive from outside the project itself. AI does not just challenge this system; it is incompatible with it.

The micro-firm structural condition and deep fragmentation of the industry prevent the broad data accumulation that AI requires (Laberge, 2025). Furthermore, project-based transience prevents the long-term, sustained collaboration that a systemic digital innovation like AI demands (Samuelson & Stehn, 2023, p. 388). At the organisational level, the strict survival logic governing these smaller builders prevents the accumulation of slack resources and the long-term focus that technology adoption necessitates (Meng & Brown, 2018, p. 1215). Simultaneously, regulatory heterogeneity prevents the repeatability that scaling requires (RBC, 2024), and financial misalignment prevents the capital flow that modular and AI-enabled innovation relies upon. (Kurup et al., 2024; RBC, 2024) Finally, the structural vacancy of a central systems integrator prevents the unifying function that systemic transformation demands. Critically, these conditions do not operate in parallel – they cascade, with each barrier reinforcing the next and closing off the adaptive responses that might otherwise allow the system to absorb innovation.

Each condition not only compounds the problem – it actively obstructs a solution that might otherwise compensate for another condition's effects. That is what makes the system *Architecturally Hostile* for AI adoption, rather than merely difficult. It is this architectural hostility – not any single barrier in isolation – that present theoretical frameworks have failed to adequately explain, and that this research addresses.

2.4.2 The Theoretical Gap: What Slaughter and Rogers Cannot See

The combined failure of Slaughter and Rogers points to a single, critical blind spot: neither framework possesses the analytical tools to interrogate the commercialisation interface. Slaughter looks at the innovation, and Rogers looks at the adopter, but standard diffusion frameworks consistently fail when applied to complex, networked technologies because they ignore the volatile places where these actors actually meet (Lyytinen & Damsgaard, 2001, pp. 7-8). Neither framework considers the space between discovery and implementation, where transactions happen and where innovations ultimately survive or die. That space, which is widely recognised as the "Valley of Death" in construction technology, has been identified empirically as a severe structural and institutional gap, yet it has never been adequately theorised within this specific industry context. (Markham, 2002, p. 31) The absence of this integrated analytical lens means that construction technology research has repeatedly diagnosed symptoms – such as adoption failure, trust erosion, and diffusion lag – without being able to explain the underlying mechanisms producing them. While recent scholarship has recognised the necessity of meso-level analysis to untangle these institutional complexities (Kussl & Wald, 2025, p. 382; Papadonikolaki et al., 2020), none have theorised the specific dynamics of the commercialisation interface with sufficient precision to account for innovator characteristics, structural vacancies, and logic clashes operating simultaneously in the Canadian residential context. This collision is magnified by the bifurcated market, which complicates the interface further. It is exactly here, at this commercialisation interface, that the theoretical silence is loudest and an adequate analytical framework is most urgently needed.

2.4.3 The Unit of Analysis: The Interaction Interface

would miss the interface. Isolating the project means the commercialisation dynamic is unaccounted for. Looking at the macro-level industry misses the micro-level transaction where adoption actually succeeds or fails. The meso-level interaction interface captures all three dimensions at once, which is what the theoretical gap requires. The work of Star and Griesemer (1989) on "Boundary Objects”, analytical constructs that sit at the intersection of different social worlds, functions similarly to the interaction interface (Star & Griesemer, 1989, p.393). Boundary objects are flexible enough to adapt to local needs, but rigid enough to maintain a common identity across groups. The interaction interface is a specific instantiation of boundary object theory applied to the novel context where startups and builders must find enough common ground to transact, without either group abandoning their institutional obligations. However, boundary objects are typically stable artefacts that mediate between communities. The interaction interface differs as it is a dynamic, contested space rather than a stable object – it is produced and reproduced through each commercialisation encounter rather than existing independently of them. And crucially, it can fail entirely if the credibility threshold is not cleared, which is not a condition Star and Griesemer's framework accounts for.

The commercialisation of AI in the built environment requires a critical analytical reframe from competition-driven to interaction-driven models. Traditionally, construction's procurement logic is competition-driven: the lowest bid wins, suppliers compete fiercely for contracts, and information is strategically withheld to maintain market leverage. Conversely, AI commercialisation relies on an interaction-driven logic that demands collaborative co-creation, transparent data sharing, and sustained, long-term relationships. At the commercialisation interface, these two logics are institutionally incompatible, requiring the active negotiation of shared meaning across institutional boundaries (Star & Griesemer, 1989, p. 393). Building on

Havenvid's core framework, Kussl and Wald (2025, p. 383) provide the clearest articulation of this *Competition Versus Interaction Paradox* in the existing construction literature, noting how simultaneous implementation of these opposing models creates severe tensions. To investigate exactly how this paradox unfolds empirically, this study targets the meso-level interface as the primary unit of analysis.

Ultimately, meso-level analysis is required to examine the interface. Macro-level analysis – such as industry trends, policy frameworks, and aggregate productivity data – cannot explain why two firms operating within identical structural conditions produce vastly different adoption outcomes. Conversely, micro-level analysis – which focuses on individual technology acceptance, user behaviour, and cognitive adoption models – cannot explain why structurally capable firms systematically fail to adopt innovations at the commercialisation stage. The meso-level is the only analytical register capable of holding structural conditions and transactional dynamics in simultaneous view, while also accounting for intermediate institutional actors (Leonardi & Barley, 2010, p.36) . Explaining commercialisation failure requires this nuance. Targeting this specific unit of analysis is methodologically justified when investigating complex, contemporary phenomena in their real-life contexts, particularly when the boundaries between the phenomenon and the surrounding organisational context are not clearly evident and require an in-depth, multi-faceted exploration (Yin, 2018). Investigating this interface requires an analytical toolkit capable of simultaneously addressing institutional logic clashes, organisational paradoxes, cognitive misalignments, and the entanglement of social and material practices – four dimensions that no single existing framework addresses alone

2.4.4 The Analytical Toolkit: Four Lenses and Why These Four

Institutional Logics, introduced by Friedland and Alford (1991) as a set of material practices and symbolic constructions that constitute its organising principles (p. 246) and operationalised by Thornton and Ocasio (2008, p. 119) for empirical organisational research, provides the analytical foundation for examining why friction at the interface is not just a commercial but institutional. When applied to the commercialisation interface, the framework analyses how startups and builders respond to the institutional complexity created by the clash of competing logics of digitalisation, project delivery, and firm survival. Critically, this lens will be used to explain why friction at the interface is not only about price or features, but that actors on either side operate from different socially constructed systems of meaning, value, and authority. Institutional Logics is the only framework that can explain why the same digital tool is simultaneously seen as transformative by one party and threatening by another.

While Institutional Logics explains why the clash exists, it cannot explain why firms fail to resolve it even when the incompatibility is clearly recognised – a tension better understood through the lens of Paradox Theory.

Paradox Theory, defined by Schad et al. as a "persistent contradiction between interdependent elements" (2016, p.5), was developed from the basic framework of Smith and Lewis (2011). This lens seeks to understand how organisations can simultaneously address competing demands and views these paradoxes as problems to be managed, not solved. When applied to the interaction interface, paradox theory can explain the persistent, irresolvable tensions that Digital Transformation (DT) generates within construction firms. This lens is necessary because it provides the analytical vocabulary for understanding why innovation adoption fails to diffuse even when it is clearly beneficial - the choice itself generates contradictions that threaten firm survival.

While Paradox Theory explains why construction firms struggle to act on innovation even when its benefits are recognised, it cannot account for why the startup and the builder regularly fail to recognise the same problem at all – a cognitive misalignment better explained through the lens of Technological Frames.

Technological Frames developed by Orlikowski and Gash (1994, p.176) specifically to examine how different organisational groups interpret the same technology differently. Central to their argument is the implication that when different groups within the same organisations understand the same technology differently, this leads to the failed implementation of new technologies. This lens will be used to explain the cognitive misalignment between IT and construction logic at the interface, or to explain why the same AI tool can be perceived to have different purposes, risks, and value propositions. Technological frames explain why this cognitive divergence is systematic rather than accidental and why innovators with domain expertise can partially bridge it.

While Technological Frames explains why cognitive misalignment at the interface is systematic rather than accidental, it cannot account for how AI tools – once adopted – foundationally reshape the authority structures and professional identities of construction management, a dynamic captured by the lens of Sociomateriality.

Sociomateriality, developed by Orlikowski (2007, p. 1437), provides the final and perhaps most humanly consequential lens in this analytical toolkit. Where Technological Frames explains cognitive misalignment at the interface, Sociomateriality explains what happens when that misalignment reaches the site level – where AI's algorithmic authority confronts the embodied, experiential knowledge of PM’s and site superintendents. Orlikowski's core premise of sociomateriality is that "there is no social that is not also material, and no material that is not also social". Positing that digital tools do not only assist humans in working practices, but instead reshape social relations, professional identities, and authority structures. (Orlikowski, 2007, p. 1437). In construction, where authority is legitimised by seniority, physical presence, and accumulated site experience rather than data analytics (Saqib & Allen, 2024, p. 2849), AI tools are not experienced as productivity improvements but as epistemological threats – systems that know differently and therefore implicitly delegitimise the way experienced practitioners know. This lens is essential for explaining why the “tried and true” versus potentially better tension observed at the interaction interface is not irrational conservatism, but a rational defence of the professional identity construction authority is founded.

Actor Network Theory (ANT), Transaction Cost Economics (TCE), and Resource Dependence Theory (RDT) are all plausible candidates for studying innovation adoption in fragmented industries.

ANT, developed by Law (1992) and Latour (2005), shows how innovations travel through layout networks of human and non-human actors by deliberately suspending questions of institutional value and social meaning in favour of following associations. Bruno Latour (2005, p. 10) expands on ANT by formulating it as a strict methodological rule: researchers must treat non-humans as active "actors" that do things, not just as tools used by humans. ANT has no mechanism for the institutional logic clashes and credibility thresholds that determine whether those networks form at all, making it insufficient for examining the specific dynamics of this research gap. Where ANT fails to account for the institutional conditions that precede network formation, Transaction Cost Economics offers a different but equally partial account – one that explains governance choices under uncertainty while reducing actors to rational calculators incapable of the identity-driven resistance that your research documents.

Williamson (1994) explains that TCE is the "cost of running the economic system," and his framework relies on behavioural assumptions about actors: *bounded rationality* – the recognition that humans intend to act rationally but are cognitively limited—and *opportunism*, the tendency to act in self-interest (Williamson, 1994, pp. 101–102). While TCE explains why firms make certain governance choices under conditions of uncertainty, it treats actors as utility maximisers – a behavioural assumption that cannot account for the institutional values and professional identities that make construction managers resist AI even when the return on investment (ROI) case is clear. While TCE's limitations lie in its behavioural assumptions, Resource Dependence Theory offers a different but equally partial account—one that captures power asymmetries at the interface without reaching the cognitive and institutional dimensions that determine its outcomes.

Pfeffer and Salancik (1978, p. 3) introduce RDT as an open-systems perspective, positing that organisations are not completely autonomous entities. Instead, organisational behaviour is constrained by the demands and pressures of the external groups upon which the firm relies for critical resources. RDT usefully explains how organisations manage these environmental dependencies and how power asymmetries dictate inter-organisational relationships (Pfeffer & Salancik, 1978, p. 108). While this macro-structural perspective is highly relevant for understanding why micro-firms in the construction industry are structurally vulnerable to technology vendors, RDT lacks the analytical tools required to investigate the cognitive misalignments, institutional logic clashes, or sociomaterial authority threats central to understanding why the commercialisation interface fails before structural power asymmetries even become relevant.

Where each of these alternatives captures one dimension of the commercialisation interface, none can simultaneously address the institutional incompatibilities, organisational paradoxes, cognitive misalignments, and sociomaterial authority threats that together explain why AI adoption fails in Canadian residential construction. Together, these four frameworks constitute an analytical toolkit whose combined coverage matches this multi-dimensional complexity – with Institutional Logics identifying the source of interface hostility, Paradox Theory explaining its persistence, Technological Frames mapping its cognitive geography, and Sociomateriality revealing its human cost. Each lens picks up precisely where the previous one reaches its analytical limit, producing sequential coverage from institutional structure to individual professional identity that the complexity of this research gap demands. With this analytical foundation established, the following sections apply these lenses to the specific dynamics of the interaction interface, beginning with the structural void that sits at its core.

2.4.5 The Structural Void: Scale, Vacancy, and the Bifurcated Market

The barriers identified across the preceding sections do not operate as isolated structural conditions. They converge into a unified structural void whose three defining dimensions collectively prevent the conditions necessary for AI commercialisation to succeed. The scale problem has been established; micro-firms lack slack capital for adoption, but they also lack the organisational bandwidth to participate meaningfully in the co-creation process that AI commercialisation requires (Laberge, 2025; Zachariah & Goldsmith, 2022). The structural vacancy of a permanent integrating actor, driven by project-based discontinuity, compounds this bandwidth deficit further (Samuelson & Stehn, 2023, pp. 388-389). The startup is forced to perform the integration function itself. Startups attempt to bridge the logic clash because they need the incumbents, which leads to overselling solutions and an erosion of trust (Bridgit, 2023). This commercial pressure to perform the integrating function while simultaneously selling a product creates a structural conflict of interest at the interface. The innovator cannot objectively assess fit when their survival depends on securing adoption. Moreover, the startups must manage the data infrastructure and stabilise the adoption process across project boundaries, compensating for the missing integrator (Samuelson & Stehn, 2023, P. 389).

This structural void means domain expertise is a functional necessity. The domain-expert innovator is partially filling a structural role the system cannot generate internally – both measuring and attempting to pass the credibility threshold. This means that the threshold is also dependent on whether the innovator has the competence and bandwidth to perform the integrating function to fill the void. The innovator, therefore, faces a structurally impossible task, but this is multiplied by the tension of a bifurcated market (Kussl & Wald, 2025, P. 383).

The structural void produces two distinct failure modes simultaneously, reflecting bifurcated innovation behaviour: micro-firms fail to adopt because a survival logic forecloses the bandwidth and capital required, whereas large firms fail to adopt because analysis paralysis and organisational inertia prevent commitment (Meng & Brown, 2018, pp. 1217-1218). The scale problem produces bandwidth poverty, bandwidth poverty prevents co-creation, the absence of co-creation means innovations arrive at the interface without incumbent input, and the missing integrator means no stabilising actor exists to compensate. This creates a self-reinforcing cycle that the bifurcated market renders doubly intractable by ensuring that no single intervention can address both failure modes simultaneously (Kussl & Wald, 2025).

Innovators face a market where neither segment can adopt their innovation through conventional commercialisation, which is why the interaction interface requires a different logic than standard technology diffusion models assume. It is within this structural void – characterised by scale insufficiency, integrator absence, and market bifurcation – that the institutional logic clash at the commercialisation interface becomes not merely difficult to navigate but structurally irresolvable without deliberate framework intervention.

2.4.6 The Clash of Logics: Product vs. Project

The friction at the commercialisation interface is not primarily commercial disagreement created by institutional differences. The friction is a direct consequence of actors operating from logics that were never designed to encounter each other. Applying an institutional logics framework (Thornton & Ocasio, 2008), we see that the startup and the builder not only disagree on price or features; they navigate different systems of meaning and authority. The startup is driven by a Logic of Digitalisation, pursuing what Schildt (2022, p. 235) describes as "digital omniscience" and "digital omnipotence." In practice, this manifests as a sales pitch built around data dashboards, interoperability, and long-term ROI projections. Conversely, the construction firm is governed by a Project Logic (Samuelson & Stehn, 2023, p. 388), where value is defined by the successful delivery of a physical asset within a rigid "one-off" timeframe. Here, the logic manifests as the immediate, pragmatic question: *"Will this tool slow down my current job?"* For micro-firms, this is further complicated by a Survival Logic (Meng & Brown, 2018, p.1215). For these actors, engagement itself has a cost their margins cannot afford, manifesting as the decision to avoid the meeting entirely – a phenomenon of *Operational Capture*. Finally, a Corporate Logic (Saqib & Allen, 2024) within larger firms prioritises hierarchy and immediate revenue, effectively suppressing bottom-up innovation by keeping digitally capable staff tethered to routine billable tasks rather than transformational implementation.

When these logics collide during the commercialisation encounter, the result is a systemic failure of interpretation. The startup's Logic of Digitalisation frames the conversation around scale and algorithmic optimisation, while the builder's Project Logic frames the same interaction around immediate delivery risk and proven reliability. These are not just different perspectives; they are categorical standards of value that cannot be fixed in a single transaction. This collision is exacerbated by incompatible survival imperatives: the startup's survival depends on securing rapid adoption within a 12-to-24-month iteration cycle, while the builder’s survival depends on ensuring that no new variable disrupts the current project’s delivery timeline. Because these actors operate on different timelines and risk tolerances, they inevitably talk past each other. The encounter fails not because of a lack of incentive, but because the institutional logics they carry into the room produce entirely different interpretations of the same technical information.

This institutional hostility is formalised in the Competition Versus Interaction Paradox, drawing on the work of Havenvid (2015) and Kussl and Wald (2025, p. 383). Construction’s dominant procurement logic enforces fierce competition where the lowest bid wins, and information is strategically withheld to protect margins. However, the commercialisation of AI requires deep interaction, collaborative co-creation, transparent data sharing, and sustained trust. These two institutional requirements are mutually exclusive at the interface; a builder cannot compete aggressively for a low-margin contract while also collaborating openly with an AI vendor whose tool requires access to operational data. This collision of institutional logics, which operates within a structurally hostile void, mediated by an analytically neglected interface, and exacerbated by a technology that refuses to stabilise, is what this research addresses, and that existing frameworks have proven inadequate to explain.

2.4.7 Thesis Contribution: From Adoption to Logic Reconciliation

The primary contribution of this research is a reframe that shifts the focus from traditional technology adoption models to the meso-level relationship-building, or mending of clashing institutional logics at the commercialisation interface. By moving away from a view of static technology or cultural factors, the study prioritises the characteristics of the innovator and the structural dynamics of the encounter itself. This shift redefines the commercialisation of AI not as a transactional hurdle, but as a complex institutional negotiation within a previously neglected analytical space.

Theoretically, this research establishes the interaction interface as a dynamic, threshold-dependent site where structural vacancies, logic clashes, and innovator characteristics collide simultaneously. Central to this theoretical contribution is the *Credibility Threshold* – identified here as a prerequisite condition for interface engagement rather than a mere competitive advantage. Without meeting this threshold, the interface remains partially broken, as incumbents are unable to evaluate technical innovation independently and instead use the innovator's domain expertise as a benchmark to judge an innovation's fit.

Building upon these theoretical mechanics, the research provides a practical contribution through a conceptual framework designed to guide startups, incumbents, and policymakers through the industry's structural void. This framework identifies specific systemic failure modes, such as the *Moving Target Trap* and *Operational Capture*. This aims to provide stakeholders with the diagnostic tools necessary to move beyond trial-and-error strategies. By operationalising these concepts, the thesis offers a roadmap for navigating the hostile interface, providing the explanatory power that existing frameworks have proven inadequate to deliver.

3 Methodology

3.1 Research Design: Exploratory Multi-Case Study

This research adopts a qualitative, exploratory, multi-case study design. This approach is best applied to explanatory questions in complex, real-world environments like construction project management. By focusing on multiple cases (the interactions between different start-ups and construction firms), the study aims to identify patterns in barriers and enablers that a single-case study might miss. The unit of analysis is the "Interaction Interface" between the digital innovator (Group A) and the industry incumbent (Group B).

3.2 Sampling Strategy: Purposeful Selection

To reach thematic code saturation, this study targets a minimum of 12 participants. This sample size is consistent with qualitative saturation thresholds identified in prior methodological research. Participants are selected using criterion-based, purposeful sampling to ensure they possess "rich information" regarding the phenomenon:

Group A (Digital Innovators): Founders or product leads from Canadian AI start-ups, ConTech consultants with experience in go-to-market strategies, and GC Free agents (Skilled innovators within the Public Service of Canada who move between departments to work on specific, high-impact projects.)

Group B (Industry Incumbents): Project managers (PM’s) or executives from residential construction firms.

Selection Markers:

Group A: Focus is placed on participants working in teams/on projects involving Agentic AI solutions.

Group B: Focus is placed on "Early Adopters" and "Sceptics" to capture the full spectrum of the diffusion lag. Furthermore, a focus is placed on participants across all areas of residential construction (single-family domiciles, Multi-family units, Large and complex buildings) to ensure findings capture industry-wide data.

3.3 Data Collection: Semi-Structured Interviews

Data is gathered through semi-structured interviews, which provide a roadmap rather than a script. This format allows the researcher to build a rapport necessary to probe deeper into unexpected insights while ensuring that core topics are covered consistently across all 12 sessions. The interviews are approximately 45 minutes per participant. Hosted on video calls that are recorded so that a verbatim transcription can be obtained.

3.4 Data Analysis: Inductive Thematic Analysis

The collected data is analysed using a systematic six-step thematic analysis process:

- **Familiarisation**: The process begins with an immersion in the interview transcripts, during which initial memos are recorded to capture preliminary insights.
- **Initial Coding**: In Vivo coding is employed to identify recurring concepts using the participants' own terminology, such as "data silos" or "margin pressure".
- **Generating Themes**: Codes are subsequently grouped into broader analytic categories to identify overarching patterns.
- **Reviewing Themes**: Themes are cross-referenced against the original primary data to ensure accuracy and to maintain objectively comprehensible results.
- **Defining Themes**: The specific "micro-level dynamics" discovered during the analysis are named and refined to ensure they address the research question.
- **Writing Up**: The final stage involves weaving the findings into a connected narrative that links the empirical data to the research aims.

This research utilises Sim-to-Real Transferability as a primary analytic lens for framework evaluation. The term originates in robotics and engineering, where it describes the gap between how a system performs under controlled simulation conditions and how it performs in variable, messy real-world deployment. In this research context, it is applied analogically. The conceptual framework derived from qualitative interview data represents a controlled theoretical environment, while the structural realities of Canadian residential construction represent the messy, institutionally complex real world the framework must ultimately address.

The Sim-to-Real lens is particularly appropriate for evaluating AI commercialisation frameworks because AI tools themselves face a messy transferability challenge. AI solutions that perform well in controlled development environments frequently fail when deployed in the variable site conditions that characterise residential construction. Applying the same evaluative logic to the framework that the framework applies to the innovations it addresses creates an internally consistent evaluative methodology.

The Sim-to-Real lens asks three structured questions of the framework. Does it capture the relevant variable: Are all documented structural barriers addressed without overclaiming beyond the empirical evidence? Does it make realistic assumptions about actor behaviour: or does it require behaviour that the structural conditions documented in the findings make unlikely? And does it account for environmental variability: does it hold under different market conditions, firm sizes, and regulatory environments? These three questions structure the framework evaluation presented in Chapter 6, where four literature-based stress tests are applied to identify the specific conditions under which the framework's transferability is strongest and weakest.

3.5 Researcher Positionality, Reflexivity, and Bias Handling

Qualitative research recognises the researcher as the primary instrument of data collection and analysis, so the researcher’s background, values, and experiences inevitably shape the "Interaction Interface" under study. This section outlines the reflexive stance adopted to ensure the integrity and trustworthiness of the findings.

3.5.1 Researcher Positionality

The researcher approaches this study with an "Insider-Outsider" duality:

With a formal background in digital transformation, the researcher possesses a high degree of "technical literacy" regarding Generative and Agentic AI. This experience facilitated deep probing during interviews with Group A (Innovators) which allowed for a technical evaluation of AI "Sim-to-Real" transferability. Conversely, the researcher possesses minimal experience in residential construction. While this represents a lack of domain expertise, it functioned as a reflexive advantage during interviews with Group B (Incumbents). The researcher occupied a "naive" position, which forced participants to articulate "taken-for-granted" industry norms and survival logics that an industry veteran might have overlooked or left unsaid.

3.5.2 Reflexivity and the Pro-Innovation Bias

A significant risk in digital entrepreneurship research is the "pro-innovation bias". The unconscious assumption that technology adoption is inherently beneficial and that resistance is irrational. Throughout the data collection process, the researcher used the participant sampling combined with bilateral confirmation as strategies to monitor personal leanings toward the "Logic of Digitalization". The researcher intentionally interviewed sceptics within the incumbent sampling group to ensure the "Hostile Interface" was documented from a defensive, rather than just an optimistic, perspective.

4 Empirical Findings

"Nothing about building houses has changed in thirty years. Some people still track projects in written notebooks." (P4, p.316) This observation, offered by a senior incumbent PM, captures the empirical reality that this chapter documents and analyses – an industry whose operational practices have remained largely frozen while the technological landscape around it has transformed beyond recognition. This section presents the empirical findings from 12 semi-structured interviews conducted with innovators, industry incumbents, and dual-perspective participants who operate in Canada's residential construction commercialisation interface. Rather than summarising what participants said, this chapter analyses what their collective accounts reveal about the structural and institutional dynamics of the interface. Using direct quotes from participants to support claims, where relevant, findings are connected to the theoretical framework established in Chapter 2, confirming, extending, or complicating existing literature.

This chapter is organised into four thematic sections that emerged through the inductive coding of the data. Each theme will address a distinct dimension of the commercialisation interface, which, when viewed together, will present a multi-dimensional empirical picture of the interface as experienced by the actors navigating it. The four themes appear as follows: The Frozen Industry, The Hostile Interface, Structural Incapacity, and Conditional Optimism.

The participants interviewed are categorised into three types and identified throughout using coded designations: Innovators(I1-5), Professionals (P1-5), and Dual-perspective participants (D1-2). Bilateral confirmation of a theme or subtheme between groups will be treated as the strongest form of evidence. While disagreement or contradictory perspectives will be examined carefully and juxtaposed to take a moderate approach. No one perspective will be treated as objectively primary to another, though backgrounds and experience will be taken into consideration, particularly when it is applied to one side of the interface.

4.1 The Frozen Industry

4.1.1 Thirty Years of Stagnation

Stagnation in Canadian residential construction is not experienced as a crisis requiring resolution – it is experienced as the natural state of things. Interviewees on both sides of the interface are aware of the productivity gap, as one innovator participant stated, "You can look at the stats and construction on productivity, and they're pretty alarming. Like not many other sectors saw a precipitous decline in the last five years and kind of a secular stagnation for the last like, I don't know, 30 (years), give or take."(I4, p. 284), while another said, " it's typically older. guys, men who are in the space and, you know, just aren't interested in changing. They, you know, do 5, 6 projects a year and want to keep it that way, you know, don't really care about, you know, optimising that sort of thing. (I1, p. 9). That both observations came independently from innovators suggests a shared perception across the startup community that incumbent satisfaction with stagnation is a structural feature of the market rather than an individual attitude. Illustrating this point, one incumbent PM interviewed for this study still actively tracks projects using written notebooks – not as an anomaly, but as standard practice within their firm – illustrates how thoroughly this stagnation has been normalised at the operational level. This pattern is particularly pronounced among micro-firms, where building two to three projects a year represents a self-contained and financially viable business model that innovation would disrupt without obvious benefit. This is further evidenced by one incumbent, saying, "So for custom homes, like one-off custom home builders, the guy, four people or less, let's say they build two or three homes a year max. That's like, that'd be easy. That's all they do." (P1, p. 49) However, this is not an opinion shared by every incumbent, where one PM explained that the lack of productivity is not only the result of satisfied incumbents, but a shifting of expectations on builders: "... there's a long process. It's been getting longer over the years. You know, it's not getting any easier to put these projects together. There's becoming more requirements, more red tape, more code requirements, adopting new code, which we did" (P3, p.188). P3's input shows that productivity is slowed by increasing burden, not just satisfaction. Empirically, this is evidence of Operational Capture, where the builder is operationally stretched so thin that they cannot consider entering the commercialisation interface. The normalisation of stagnation as industry default – rather than crisis requiring resolution – presents a foundational barrier to AI adoption: incumbents who do not experience productivity as a problem will not come to the negotiation table, let alone participate in the co-creation of innovations. This empirical pattern confirms the secular productivity stagnation documented in Chapter 2 while extending it into the subjective dimension – participants do not solely inhabit the stagnation, they have internalised it as the natural state of their industry.

4.1.2 The Real Estate Appreciation Shield

Stagnation in Canadian residential housing is not simply a product of industry culture; it is the direct result of an environment where poor practices were allowed to persist with impunity. In the residential construction industry, companies are only spending money while they build, so the business model hinges on the profitability of the property after completion. This equation is simple: The cost per square foot to build needs to be lower than the cost per-square foot to rent or sell. Real estate appreciation, particularly in Canadian metropolitan hubs, provided a financial shield to builders, who could run poorly managed projects that routinely came in over budget, but still be profitable. This shield lessened the financial pressure on builders; this pressure is exactly what drives innovation. Over the last two decades, the absence of this financial pressure meant that the fractured market could continue the status quo, be innovation resistant, and remain profitable. This pattern represents a contextually specific Canadian barrier that extends the literature review's structural analysis – the appreciation shield operated as a market distortion that suspended the productivity pressures that normally drive innovation adoption, producing a legacy of normalised inefficiency that persists beyond the conditions that created it.

This dynamic was articulated with particular clarity by a dual-perspective participant with extensive experience in both innovation and project delivery. “There's lots of dynamics in real estate and market pricing, rents. You know, I kind of laugh. I don't know if you've ever spent time in Vancouver, but I got here in 1999 after being in Europe for eight years, and, like the price of real estate, has grown double digits for many years. Well, maybe high singles and sometimes double digits. So developers could, they were getting lazy. They just could make one mistake after another with their budgets and schedules. It didn't matter because by the time they finished the building, even if the building took longer. The market price per square foot was higher.” (D2, p. 64-65) That this observation came from a dual-perspective participant – someone with experience on both sides of the interface – lends particular evidential weight to the claim that the appreciation shield was not merely a market condition but an active suppressor of the innovation pressure that might otherwise have forced productivity improvement.

4.1.3 Adversarial Industry Structure

The frozen industry is not just technologically stagnant – it is institutionally structured for competition in ways that systematically prevent the collaboration conditions that AI commercialisation requires before the innovation encounter even begins. This industry is reliant on a complex web of subcontractors, which leads to an environment where each individual is focused on maximising their own profit and minimising their risk. In practice, this dynamic was captured precisely by a dual-perspective participant with experience on both sides of the interface "Again, it gets into the fractured nature, like…And how are you actually going to get from the plumber data that you're going to use as a general contractor? Well, the plumber basically doesn't want to give you any data. He wants to keep you guessing, right?" (D2, p. 77) The reality is, the sub-contractor loses their bargaining power when data is transparent, and views collaboration as a risk to their own margin. This every-man-for-themselves environment, combined with thin project margins, stands in opposition to progress because inter-project relations become very transactional. This was further supported by another participant who said, “There's all these different spots along the way. It's everyone's protecting their own little profit pools. There's all these handoff points. Everyone's pointing fingers at each other. So it's just a very fractious sort of environment... So if we're going to solve the problem... we've got to look at the entire supply chain.” (I3, p. 277) Participants recognised this not as a cultural failing but as a structural inevitability, as one innovator observed, “It's an industrial structure problem, right? Like, we don't have very big companies in construction... Traditional construction is basically, you know, there's a general contractor running a project, but then you have tonnes and tonnes and tonnes of subcontract subtrades. So like, that is a model that does not lend itself well to the kinds of efficiency gains that you see from centralising operations and then rationalising things because it's just not how the ownership structures work in the sector.” (I4, p. 285) Furthermore, this structural inefficiency being adversarial is felt as a constraint even by those operating within the system, as it was bilaterally confirmed by an incumbent, who explained, "So the weight is now on the contractors and everyone's very risk adverse nowadays. Don't want to be wrong, don't want to be responsible, will do whatever they can to avoid that kind of scenario of being the one caught holding the bag... It stifles the whole concept of construction and teamwork, you know." (P5, p. 249)

Taken together, these accounts confirm that the adversarial structure is not a cultural attitude that trust-building can overcome – it is a rational response to an ownership and incentive architecture that makes data sharing and collaboration structurally dangerous for individual actors. This empirical pattern directly confirms the Competition Versus Interaction Paradox identified in Chapter 2 – construction's procurement logic enforces the competitive behaviours that AI commercialisation requires participants to abandon (Kussl & Wald, 2025). With collaborative conditions thus institutionally prevented by design, it becomes inevitable that when technology startups attempt to inject digital solutions into this fractured ecosystem, they are met not with an open door, but with a profoundly hostile interface. It is within this frozen and adversarially structured industry that innovators must attempt to commercialise AI solutions – entering an interface whose hostility is not incidental but architecturally produced by the conditions Theme One has documented.

4.2 The Hostile Interface

4.2.1 The Credibility Threshold: The Grey Hair Effect

“So the reason we didn't make as much money as we should have was because…I have a project manager I started with, I brought him all the way through. He was like my first guy. I called it the *Grey Hair Effect*. As soon as I brought him on board and had him start talking to clients, no one questioned him about age or projects... trust doesn't exist unless you put your time in." (P1, p. 42) This quote came from an industry incumbent who started their own construction company, but was viewed as too young or inexperienced within the industry. Once they brought on an older PM, everything changed; they had crossed the Credibility Threshold.

The residential construction industry exists in a dearth of objective performance metrics, so incumbents use domain experience as a surrogate for trustworthiness. This was the case in an incumbent-to-incumbent interface, but when it comes to innovators, crossing this threshold becomes more difficult and more important. Another incumbent said, “You need lots of feedback from guys that have done it for 30 years, 40 years, 50 years. Grey hair means a lot in construction. We don't, I don't think we have enough. old school or boots on the ground input in this new software to really make it work well.” (P5, p. 259) This serves as double confirmation that the Grey Hair Effect is important for incumbent-to-incumbent relations, but this same incumbent also explained the importance of domain expertise on the innovator side

“You have all these, I'm going to call them dreamers just for the sake of conversation. These people that want to make software that's going to make them millions of dollars, solve everybody's problems. But they're, you know, not many 60-year-olds are doing that… I'm sure 9 times out of 10, they get 9 out of 10 things, right? But there's always that one thing that, unless you've done it before, doesn't cross your mind. And that could be, you know, the straw that breaks the camel's back, so to speak. That could be the one thing they need.” (P5, p. 258-259) This observation captures the functional dimension of the Credibility Threshold – domain expertise is not just a trust signal but a prerequisite for product validity. The innovator without construction experience cannot know what they do not know, and in construction, that unknown is often the critical variable. This pattern was independently confirmed from the innovator side of the interface, where participants acknowledged that domain expertise was a prerequisite for being taken seriously, regardless of product quality. “We started out being a problem-first company… And domain expertise is so, so important. And I think a key failure for a lot of technology companies is they develop something. It might be better, faster, cheaper, but it's looking for a problem to solve, right?” (I3, p. 273). An interesting example of how domain expertise can help cross the Credibility Threshold came from a dual-perspective participant who explained “...at the age of 31 or 32, I decided to take a two-year break and do carpentry. And that was probably a big part of my ability to be an extremely, extremely successful PM because I can have debates and challenging conversations with superintendents, foremen, and also subcontractors about how they do their work, how they organise, and what's ergonomically required, right?” (D2, p. 80) In this instance, we have an individual realising the importance of the Credibility Threshold and executing a plan for crossing it. This bilateral confirmation demonstrates that the Credibility Threshold is a structural feature of how trust operates throughout the entire industry. Innovators face a more extreme version of a universal industry condition. This finding addresses the theoretical gap identified in Chapter 2 – Slaughter's (1998) innovation typology accounts comprehensively for product characteristics but contains no variable for innovator characteristics, despite the empirical evidence presented here that innovator domain expertise is the primary determinant of whether the interface functions at all.

The Grey Hair Effect and the Credibility Threshold are rational institutional responses to a fundamental information problem – in an industry that lacks the slack resources to independently evaluate AI solutions, domain expertise becomes the only available substitute for trustworthiness. Domain expert innovators must fill this role and evaluate the AI tool – based on fit and ROI – on behalf of incumbents. This is domain expertise as cognitive offloading, not simply legitimacy signalling. Thus, the Credibility Threshold is not just a competitive advantage; it is an operating prerequisite condition.

The hostile interface documented in this section is not primarily the product of individual attitudes, commercial miscalculations, or cultural resistance. It is sustained by structural conditions that make meaningful engagement at the interface systematically difficult regardless of the participants' intentions, capabilities, or incentives. These structural conditions – operating beneath and beyond the interface itself – are the subject of Section 4.3.

4.2.2 The No First Mover Dynamic

The commercialisation interface is made more hostile because even if domain-expert innovators have genuine solutions and can cross the Credibility Threshold, they run into a social roadblock in the form of industry adoption decisions. The *No First Mover* dynamic complicates the interface further by introducing a social dimension where incumbent adoption decisions do not happen in isolation, but in relation to what peer firms are doing. Not only was this dynamic independently verified by interviewees on both sides of the interface, but it also connects to the literature of herd adoption introduced in Chapter 2 (Lyytinen and Damsgaard, 2001). One incumbent described this as "no one's going to want to be the first person in the guinea pig to try out these systems" (P1, p. 52) in reference to new software or AI solutions. Then confirmed by a dual-perspective participant “even if your price is the same with a new solution... risking that, even if the price is right, it's the biggest bottleneck the that being the first one... nobody wants to be the first one to try something new.” (D1, p. 130). Again, this dynamic represents a rational response to structural industry conditions rather than reflexive conservatism or resistance to change. First-movers bear all the risk of failed implementation or faulty technology while their peers watch and learn. The firm that adopts second is taking on a risk that has been partially socialised and inherently protects itself from dire consequences. While this rational behaviour is justifiable on an individual firm level, it results in industry-wide irrational outcomes where effective innovations fail to be adopted because no firm moves first. This is a classic failure in coordination that can be explained by the Competition vs Interaction Paradox introduced in Chapter 2.

Furthermore, incumbents are not just afraid of being first; they are also waiting for the new technologies to stabilise before committing. This is the *Moving Target Trap* impacting individual decision-making. AI technologies are evolving rapidly, and risk-averse incumbents take a “wait-and-see approach” to AI tools. Empirical evidence of this was reported by two independent innovators. The first, in reference to cold outreach with incumbents (actual interaction at the interface), “... some were just like, this sounds really cool, but you know, we want to wait till the software gets a little more reliable, which is, you know, fair. Like it does take time out of their day. You know, if it's wrong, it's going to cost them, it's going to, you know, be a headache even more than it is a help” (I1, p. 11). Innovators clearly understand the friction and can make the distinction between irrational laggards and rational conservatism. The second innovator explained a similar instance, where incumbents showed interest, but still opted to wait-and-see “There are a lot of people that just say, yeah, interesting what you're doing, but how am I going to make more money?... tell me when you're ready to go and I'll decide then.” (I3, p. 281) This deferral logic, while individually rational, empirically confirms the Moving Target Trap – the stability being waited for is structurally unattainable within AI's current developmental trajectory, meaning the decision to wait is functionally equivalent to a decision not to adopt.

The last factor contributing to the no first mover dynamic is the bandwidth cost variant. A dual-perspective participant explained the thought process of incumbents, “Well, it's working what I got. I'm making money. I'm doing projects. Long comes a new technology, shiny. How do we know it's actually going to work? And how much time am I going to get consumed, am I going to consume into having people adopt it?” (D2, p. 86) This is Operational Capture again, only this time as a specific barrier to first-mover adoption. This extends the risk of adoption decision-making past just reputational cost, but onto the change management that would disrupt project delivery - a disruption most firms cannot absorb.

The structural nature of this coordination failure points toward an intervention logic that at least one participant independently identified. One dual-perspective participant identified a logic that this thesis's framework will develop – that the no-first-mover coordination failure requires an external actor to socialise adoption risk through funded pilot programmes before voluntary first-mover behaviour can be expected. The participant noted, “So CMHC just started, like I think they funded a couple of pilot projects, so pilot projects I think it's a great idea or incentive of giving new technologies a chance of being used uh and demonstrate and then having an actual real use case." (D1, p. 130)

4.2.3 The Hallucination Trust Deficit

AI's specific hallucination errors introduce a trust deficit that standard technology adoption frameworks have no mechanism to address – unlike previous construction technologies whose limitations were knowable and bounded, AI errors are unpredictable, hidden until discovered, and consequential in a high-stakes physical delivery environment One Innovator explained that while their solution is using AI, the core of their solution is automation, this decision was made because "You know, AI hallucinates, there's, you know, it's not always trustworthy. It can, you know, make some kind of uncalled for decisions that, you know, might not be information backed, because that's just how these algorithms work through, you know, piecing together information from, you know, large, large sources of data…" (I1, p. 4) Clearly, AI hallucinations are taken seriously enough by innovators that one recognises the effects they can have on incumbent trust. A dual-perspective participant acknowledged hallucinations, but provided a more nuanced view of specific functions where hallucinations are more troubling than others "ChatGPT was hallucinating completely, like not making sense... I'm not that worried about hallucination because I feel that models are getting better with Excel sheets, and when it comes to math, just simple, straight up math. I think that they tend to hallucinate more on just not numerical inputs uh and just interpretation of ideas or concepts. That's where I see more hallucination." (D1, p. 129). This adds a layer to this analysis as some innovators are more optimistic than others regarding AI hallucinations. This optimism was corroborated by an innovator who explained, "...my understanding and like my own usage is that the hallucination rate has, I think, gone down. So, you know, it's I think always going to be an issue and you're always going to want a human in the loop to, like, make sure that things are accurate, especially when there's a lot of money on the line. But yeah... an issue that I would say I am decreasingly concerned about." (I3, p. 288) On the innovator side of the interface, it is clear opinions are divided, and this spectrum of opinions compounds the trust rift between innovators and incumbents by sending inconsistent signals to construction firms. The emergence of hallucination rate as a participant-generated metric for evaluating AI trustworthiness, and human-in-the-loop as a way to mitigate risk from hallucinations, suggests that incumbents and innovators are beginning to develop shared evaluative language for AI reliability – a precondition for more productive interface engagement. On the incumbent side of the interface, even when probed, hallucinations did not come up as a major concern in any interviews, but more as a generalised concern about AI-generated information. One participant noted difficulty distinguishing trustworthy from untrustworthy AI outputs in daily information consumption. (P2, p. 158)

AI hallucinations are errors, and just like any technology, there is going to be a tolerance for errors. A dual-perspective participant mentioned that they would accept an AI-delay-prediction-tool based on an error rate “I would accept a statistical divergence. amount of error, like hallucination. ... If it had 20% was incorrect, I would be okay with that. Maybe 30%, 30%” (D2, p. 93) The participant acknowledged that this is would only apply to an environment where the AI tool is being used by an employee, and outputs would be checked each time, explaining ‘it's not the idea that we would completely offload all of the cognitive work to this tool and trust it completely, but that whoever's job it is to foresee errors or a project, let's call it. delays or whatever, if they have this tool that gives them information, maybe, yeah,” (94). This tolerance for statistical divergence suggests that adoption barriers are not uniformly about reliability thresholds – at least some incumbents have developed a pragmatic error tolerance that treats AI as a probabilistic tool rather than a deterministic one. Whether this represents sophisticated risk calibration or underestimation of consequential error costs remains an open question that the data cannot resolve.

The reliability threshold that incumbents require is a moving target. Innovators consider the problem to be eventually ameliorated by future technological improvement. This pattern extends the Moving Target Trap from the strategic level identified in Chapter 2 to the product level – the specific reliability threshold that would satisfy incumbent risk tolerance is perpetually deferred to the next iteration of AI development, making it structurally unattainable within any given adoption decision window.

One participant offered an empirically tested technical mitigation for the hallucination trust deficit: restricting AI models to closed-loop domain-specific data pools using retrieval-augmented generation methodology. By providing the AI only with certified zoning information and pre-configured design parameters, hallucination risk was effectively eliminated within that operational scope. (I5, p. 305) This suggests that the hallucination trust deficit, while real, is not permanently irresolvable. Potentially, it is addressable through deliberate data architecture decisions that constrain AI outputs to verified domain knowledge

4.2.4 Overselling and Trust Erosion

The commercialisation interface is systematically distorted by a structural conflict of interest that makes honest product representation institutionally impossible for start-ups. These start-ups are operating on timelines that are completely incompatible with construction's adoption pace. Start-ups need adoption to gain funding, and then use it to grow and scale. Whereas construction companies need profit to survive. This mismatch inevitably means that startups will oversell the capabilities of their solution. This leads to products that do not live up to expectations, which erodes trust between the industry and the innovator. This dynamic was captured by a dual-perspective participant who demanded evidence of fit that the startup could not provide: "I don't want a demo, I don't want to spend half-an-hour chit-chatting about how great you think your software is. Stop, let me look at your company, do some research, I'll send you an Request for proposal (RFP) document. And then I send it to them, and they're like... Ohh. Well, we can do some of that. I said, well, how much of some of that can you do? Well, you know, and they just, they're almost like, almost resisting trying to answer the RFP because they know that they can't" (D2, p. 87) Coming from a dual-perspective participant, this is particularly illuminating to the kind of tension that exists at the interface when the incumbent is demanding evidence of product fit that the startup cannot provide. In this instance, the RFP is a reasonable request of evidence, and the startup's inability to answer it fully is a structural feature, not a commercial ethics failure. This erosion of trust is cumulative and industry-wide – each failed promise raises the evidential threshold for the next innovator, regardless of that innovator's product quality. This dynamic was explained by I4, a GC free agent, as "Oftentimes you have people who might have really great solutions, but I think also, you know, there's a history of people feeling burned by, you know, things that didn't pan out, didn't live up to their promises. And like I said, like most people in the construction industry, from what I've experienced, are kind of too pressed on their day-to-day to like... worry too much about, you know, looking past, you know, oh, if we get through the learning curve on this, it It could be really powerful." (I4, p. 292*)* This tech fatigue is cumulative; each failed promise raises the evidential threshold for the next innovator. This quote from I4 mirrors almost exactly what Chapter 2 reported on the warnings outlined by Bridgit (2023).

A dual-perspective participant plainly spoke of the genuine shock they felt at discovering that a startup was openly admitting – proudly – that after 10 years, they would finally become profitable "We thought we were... We weren't doing it right. We're actually making money. But it took us like multiple years to get there. This guy is actually openly saying that we're only going to get to cash flow positive next year.” (D2, p. 113) This experience shows that startup survival logic operates on fundamentally different timescales and financial realities than construction firm survival logic. The startup CEO running at a loss for ten years is not being dishonest about their product – they're operating from an institutional logic where sustained loss is a rational investment in future market position. The construction executive making money from year one cannot conceptually process this as legitimate business behaviour. These are not just different strategies – these are different institutional definitions of what a viable business looks like. Even innovators acknowledge that AI in the tech-startup space is being commercially exploited, with one innovator explaining, “I think that's a lot of the hype. I mean, AI just is a hyped up topic right now. So I think some people are, you know, playing on that hype while others are, you know, really using it for its best purpose” (I1, p. 4) Understanding the acknowledgement of this structural pressure from the innovator side, shows that this is not a moral failing on the behalf of innovators, but instead, an opportunity for survival that even aware actors cannot fully resist.

This results in a self-reinforcing cycle. Each overselling encounter raises the Credibility Threshold for the next innovator. Higher thresholds make honest representation harder to achieve. Harder, honest representation produces more overselling. What emerges across these four sub-themes is not a collection of independent barriers but a self-reinforcing causal system – the Credibility Threshold filters out domain-inexperienced innovators, the no-first-mover dynamic stalls even credible ones, the hallucination trust deficit gives rational justification to the wait-and-see approach, and the structural conflict of interest ensures that startups respond to these barriers with overselling that raises the threshold further for the next innovator.

The hostile interface documented in this section is not primarily a product of individual attitudes or commercial failures – it is sustained by structural conditions that make meaningful engagement at the interface systematically difficult regardless of the participants' intentions. These structural conditions are the subject of Section 4.3.

4.3 Structural Incapacity

In this section, empirical evidence shows how even willing actors on both sides face structural conditions that prevent meaningful co-creation. Thus, providing proof of the Architecturally Hostile interface.

4.3.1 Operational Capture: Willingness Without Capacity

Across incumbent and innovator interviews, a recurring pattern materialised: the primary barrier to tech adoption in residential construction firms is structural capacity. Firms are genuinely interested in innovation but lack the bandwidth to translate that interest into engagement, let alone adoption. Project delivery is often an organisational-wide, Herculean effort for firms, which leaves them no slack resources for the co-creation process that adoption requires. Attributing this to risk aversion is misconstruing the issue – risk aversion is a preference, this is Operational Capture, and it is a structural condition. This is an important distinction as, in later sections, a framework must suggest interventions that untangle the actual barrier of bandwidth poverty.

When asked directly about the absorptive capacity of incumbent construction firms, one innovator's response was unambiguous in its simplicity: 'Yeah, they just can't, yeah.' (I3, p. 280) The absence of elaboration is itself analytically telling – for those operating at the interface, the incapacity is so self-evident that it requires no further explanation. This incapacity operates on multiple dimensions simultaneously. Financially, as one innovator observed, “the industry is asset light. They don't have the balance sheet. They don't adopt technologies.” (I3, p. 279) Firms without financial slack resources cannot absorb the sunk cost of a failed implementation, making even low-cost AI tools prohibitively risky when project margins are the only buffer between viability and insolvency. This empirically confirms the survival logic documented by Meng and Brown (2018) in Chapter 2, extending it from a financial constraint into a cognitive and operational one. This dynamic was independently confirmed by a dual-perspective participant who noted that even early adopters “will never have the capacity or bandwidth to really adopt” new solutions because their operational commitments prevent them from completing the implementation process. (D1, p. 136) Furthermore, Operational Capture still extends into those few firms that decide to be the first movers because change management demands are competing with active project delivery. The operational reality of this bandwidth poverty is visible at the project execution level. One dual-perspective participant described the superintendent's role as nominally including activity tracking, but noted that “they rarely have time to do that.” (D2, p. 95) If the basic operational function of progress tracking cannot be performed due to bandwidth constraints, the additional cognitive and operational load of AI tool adoption becomes structurally impossible to absorb during active project delivery.

Operational Capture does not just prevent adoption – it prevents the co-creation process that would make adoption viable and innovations valuable. Incumbents who cannot spare bandwidth to evaluate AI tools independently are structurally dependent on domain-expert innovators to perform that evaluation function on their behalf. This is the mechanism that makes the Credibility Threshold not just a trust requirement but a functional necessity – without a domain-expert innovator performing the integrating function, the bandwidth-constrained incumbent has no viable pathway to informed adoption. The distinction between unwillingness and incapacity is necessarily prescriptive. Interventions designed to overcome ideological resistance will fail against structural bandwidth poverty. Effective framework responses must address the capacity deficit rather than assuming that demonstrated ROI or trust-building alone can overcome a barrier that is fundamentally operational rather than attitudinal.

4.3.2 Data Infrastructure and Subcontractor Dependency

Data infrastructure is a prerequisite for AI adoption, but in Canadian residential construction, three structural conditions – one-off project logic, subcontractor fragmentation, and adversarial competitive dynamics – combine to make its development architecturally impossible regardless of individual firm willingness. A dual-perspective participant with extensive project delivery experience offered a striking assessment of construction data's value: “Well, it goes back to the cost, right? And material changing the material pricing, like these cost inputs, they just change all the time. I don't know how to answer that question, because the value of the data is not necessarily that great. If you're doing repetitive cookie-cutter stuff, then yeah, probably the data would be good. But the reality is we're not.” (D2, p. 75) The interviewee is speaking about one-off production rendering historically, data less structurally transferable, specifically, data that changes significantly over time. Chapter 2 flags this, and it is empirically confirmed and extended – the project-based discontinuity that Samuelson and Stehn (2023) identified at the organisational level manifests at the data level as a structural barrier to the kind of continuous learning AI requires. This data deficit is visible at the project level. One innovator noted that smaller residential projects frequently forgo 3D modelling entirely because “Why would they pay the design cost to get a Revit model for something they can do through site meetings.” (I1, p. 3) The rational per-project economics of data investment produce an irrational aggregate outcome – an industry that cannot build the data infrastructure AI requires because no individual project justifies the investment. One dual-perspective participant cut to the heart of this dynamic, arguing that the industry's real deficit is not data but collaboration – 'It's all about the data. I'm like... obviously you don't know anything about construction then. Right? It's about collaboration.' (D2, p. 76) This observation reframes the data problem as a symptom of the deeper subcontractor fragmentation that makes collaborative data collection structurally impossible. Even when firms attempt to track project data, the operational complexity of residential construction makes accurate capture structurally difficult. One innovator described how a single change order, moving an HVAC duct, can create a ripple effect across multiple interdependent systems, each managed by different subcontractors with different data practices. (I1, p. 8) The subcontractor web that produces the adversarial dynamics documented in Theme One simultaneously prevents the standardised data collection that AI tools require to function reliably. Across multiple interviews with participants on both sides of the interface, the problem of subcontractors being resistant to digital tools was prevalent. One incumbent PM explained this tension as “Some of the older guys don't have the time for it, but if you take the time, like that's the thing. If an older gentleman says, ProCore? I'll never use that. Well, you could say, all right, and you cater to them or you say, well, no, I disagree. I think you can learn it and you sit down with them and you go through it, and maybe hold their hand a little bit." (P5, p. 238) While one innovator explained the importance of a unified digitalisation effort across the whole project team: “We get the pushback... our tool works the best when all teams are using it. You know, your electrician's using it, your plumber's using it... But then, you know, we've had some general contractors say, oh, like, what if our electrician doesn't want to use it? Like, doesn't care to use it, like, what then?” (I1, p. 13) A general contractor's decision to adopt an AI tool means nothing if the subcontractors whose activities generate the most operationally critical data refuse to engage with digital tracking systems – and their willingness varies enormously and cannot be contractually compelled without damaging the relationships that project delivery depends on.

The data infrastructure problem in Canadian residential construction is not a gap waiting to be filled through better technology or stronger incentives. It is a structural inevitability produced by three mutually reinforcing conditions – the one-off project logic that limits data transferability, the subcontractor web that fragments data collection, and the adversarial competitive dynamic that makes data sharing institutionally dangerous. AI tools that require clean, continuous, integrated data ecosystems are entering an environment that is architecturally incapable of generating them.

4.3.3 Regulatory Non-Replicability Confirmed

Regulatory fragmentation in Canadian residential construction does not just slow AI adoption – it actively prevents the scaling that makes AI commercialisation viable. A solution that works in Ontario cannot be deployed in British Columbia without significant redesign or recertification. This means that startups achieving adoption in one jurisdiction are not building a scalable business – they are building a geographically locked pilot. One innovator described the geographic constraint their regulatory environment imposed on their go-to-market strategy directly: “It's very much Ontario and specifically Canada. So that's kind of early, to prove ourselves a little bit more, and then kind of going back to the drawing board when we want to expand.” (I1, p. 17) This concept was confirmed again by a GC free agent who echoed, “You know that all that rework that's necessary. So an architect or designer or construction manager, you know, they build their plans, like this is great, this meets everybody's needs. You take it to the next municipality over, and you're like… All your walls are 4 inches too short. Redesign the whole thing. Back to the drawing board. Eight months' worth of design again. Say nothing of the loss. Say nothing of the cost.” (I2, p. 219) The idea of going back to the drawing board captures the cost of expanding to a new jurisdiction. It is not a matter of scaling an existing solution but of rebuilding it to meet entirely different regulatory requirements. This empirically confirms what this thesis terms Regulatory Non-Replicability, first introduced in Chapter 2. This barrier operates on both sides of the regulatory interface. A dual-perspective participant noted that even where appetite exists to adopt new technologies, “the regulators don't have capacity or methodology to fold in these new technologies or these new innovations or materials.” (D1, p. 123) The regulatory constraint is therefore not only a compliance hurdle, but it is a structural bottleneck produced by institutional capacity limitations that no individual startup can resolve through product improvement or commercial strategy. The cumulative effect of Regulatory Non-Replicability on the Canadian ConTech ecosystem is significant. Startups that achieve adoption in one jurisdiction cannot leverage that success to build the market depth required for sustained AI development. Each new jurisdiction requires a fresh regulatory engagement, a fresh proof-of-concept, and a fresh Credibility Threshold to clear – effectively resetting the commercialisation clock and preventing the compound scaling that drives AI improvement in less fragmented sectors. This systemic misalignment was articulated by one participant: “It's very hard to produce homes that are compliant with land use and site plan regulations across a huge range of jurisdictions. If the whole issue is that we need to industrialise and we don't have a regulatory framework that allows for industrialised construction, it's a huge issue.” (I4, p. 292) This observation extends the Regulatory Non-Replicability argument beyond individual product scaling into the structural incompatibility between AI's industrialisation logic and construction's jurisdiction-by-jurisdiction regulatory architecture. However, an innovator with experience as a housing policy maker explained: “A lot of people blame municipalities, it’s not the fault of municipalities… We need life safety in buildings… I think the one of the biggest barriers to the adoption of modern construction methods is… we don’t have a standardized approval process.” (i5, p.299) This complicates matters because it shows that while the interviews empirically validates the Regulatory Non-Replicability concept introduced in Chapter 2, it does not clarify who is to blame for this issue, and how to solve it.

Resolving Regulatory Non-Replicability requires intervention at the institutional level – standardisation of AI evaluation frameworks across Canadian jurisdictions, or the creation of a national certification pathway that allows solutions validated in one province to be recognised in others. This is beyond the capacity of any individual startup or incumbent firm to achieve unilaterally.

4.3.4 The Ecosystem Coordination Failure

One innovator described the ecosystem deadlock their commercialisation strategy encountered: “It is a bit of a chicken and egg, for sure, because obviously they aren't going to do that investment until they actually see some RFPs, maybe win some RFPs. And then on the other hand, how do we know that that manufacturer is actually going to be able to produce if they aren't already producing at the scale that we need them to? So it is a bit of a chicken and egg, and we're working very closely with industry to kind of learn together and grow together.” (I2, p. 221) The value of any single innovation is contingent on the concurrent development of complementary capabilities that no individual actor controls. This is the defining condition of ecosystem failure, Adner (2006) identifies. This ecosystem coordination failure mirrors the no-first-mover dynamic documented in Theme Two – but operating at a higher level of abstraction. Where the no-first-mover dynamic describes individual firms waiting for peer adoption to de-risk an innovation before committing, the ecosystem coordination failure describes entire categories of actors. Ecosystem coordination failure occurs when manufacturers, regulators, capital providers, and construction firms are each waiting for others to de-risk the innovation environment before investing their own resources. This coordination failure extends into the capital ecosystem. Multiple participants described venture capital (VC) investment in ConTech as following herd behaviour rather than sector need. VC firms are waiting for other investors to validate a space before committing, reproducing the same coordination failure at the capital level that incumbents demonstrate at the adoption level. A dual-perspective participant explained, "That's why they're lemmings to venture capitalists... If one person decides that construction technology is now the next best thing, they all kind of run to that. And then somebody announces that we've done an investment in such and such, then they all want to do an investment in that same company." (D2, p. 114)

For one innovator, the availability of VC may be a structural problem in this industry, explaining, “we can go out and get venture capital. But venture capital wants, you know, 40% IRR. They need their money in and out, 5, 10x to return in five to seven years. They don't like investing in hardware technology, like hard technology companies, like what we're doing, because the industry does not adopt very quickly. There's a lot of capital required, which means dilution. And it takes a long time to develop this stuff. So that is access to capital is a really critical issue in this space.” (I3, p. 273) AI technology in construction could be the high-upside opportunity that VC has waited for and could result in the first VC mover.

The four structural conditions documented in this section – Operational Capture, data infrastructure failure, Regulatory Non-Replicability, and ecosystem coordination failure – are not independent barriers that could be addressed sequentially. They are mutually reinforcing dimensions of the same structural void identified in Chapter 2. Operational Capture prevents co-creation. Data infrastructure failure prevents AI tools from improving through use. Regulatory Non-Replicability prevents scaling. And the ecosystem coordination failure ensures that no individual actor has sufficient incentive to address any of these conditions unilaterally. Together, they produce a structural environment in which AI commercialisation cannot succeed through market mechanisms alone.

Yet despite these structural conditions, participants on both sides of the interface identified specific contexts in which AI adoption becomes viable, and expressed landscgenuine – if qualified – optimism about the sector's trajectory. The conditions under which the interface becomes navigable, and the futures participants imagine for AI in residential construction, are the subject of Section 4.4

4.4 Conditional Optimism

This section uses empirical evidence from participant interviews to explain the conditional optimism that is required for AI adoption in the construction industry. This theme prevents the thesis from being purely diagnostic and shows where and how AI can actually work in this industry, given everything documented in Themes One through Three: what does viable AI adoption in Canadian residential construction actually look like?

4.4.1 Where AI Actually Works: The Low-Hanging Fruit

Despite the structural barriers documented in the preceding sections, a consistent pattern emerged across interviews on both sides of the interface – participants could identify specific, narrow AI applications where adoption is not only viable but actively desired. A dual-perspective participant explained a workflow they would want AI to do: “There's no reason why Generative AI couldn't be used to prepare drawings or essentially ingest building codes and then spit out building designs based on some preliminary constraints... Why can't the review process be all done through Agentic? AI, where you are basically using human language to do checks on design.” (D2, p. 70) The interviewee is describing an AI tool that can take building codes and slightly change preset design rules to fit these codes. This is one area where AI could provide real value by reducing working-capital. However, the jurisdiction-specific nature of building codes means this application remains subject to the Regulatory Non-Replicability constraints identified in Section 4.3.3. This means AI tools can be viable within a single jurisdiction but not automatically scalable across the fragmented Canadian regulatory landscape. An incumbent PM explained that they wish AI could be used for tedious review tasks that require repetitive human intervention. Describing this ideally as, “I can throw a full set of drawings at and be able to let him to go [sic] through and review and and identify any potential issues that are going to come up during our... Fire marshal review, or the city review… parse through 1000 pages of individual trusses” (P3, p. 213). The ROI in this scenario is clear – reducing PM repetitive tasks – but this would require significant industry shifts in AI governance policies. Another incumbent, a micro-firm owner, predicted that "AI is gonna completely replace a lot of all the administration jobs” and that an estimation tool is where they want AI to work. “If you had an AI system that could do estimating for you, that would be awesome. There's no good estimating system out there that I have found yet.” (P1, p. 51) The same participant also expressed an interest in an automated change-order system on-site. This is evidence again of the willingness of the incumbent to innovate, as the incumbent has looked and has not found a suitable solution to their niche problem. Estimation was confirmed as a pain point by a dual-perspective participant who said, “I saw one that was... an estimation tool, almost like a jobber that you can like quote your own projects there... updating massive catalogues of products... being able to leverage AI to make the life of an estimator easier... it would democratise a bit the construction…” (D1, p. 137)

The potential applications of AI technology laid out by interviewees provide clear findings for AI application in this space – AI's immediate viable application is in administrative, design, and estimation tasks that are currently performed manually and are document-heavy, repetitive, and rules-based. This is the concept of *Asymmetric Adoption Threshold* – the condition whereby an innovation must simultaneously and dramatically exceed disruption cost, avoid competing with active project delivery, and feel sufficiently stable to commit to – confirmed positively. These applications clear the threshold because the disruption cost is low and the benefit is immediately visible. Broad platforms that attempt to do too much fail, but targeted workflow solutions are being actively searched. It should be highlighted that this small-scale niche application is in direct opposition to the product logic of startups, who look to maximise their addressable market. This tension between the niche solutions incumbents are actively searching for and the broad platforms startups are incentivised to build, represents perhaps the most tractable gap at the interface – one that a reoriented commercialisation strategy could address without requiring structural change. This pattern extends Rogers' (2003) compatibility concept – viable AI adoption in residential construction requires not just compatibility with existing practices but solutions so narrow and non-disruptive that the adoption process itself does not compete with active project delivery.

4.4.2 The Problem-First Imperative

Viable AI adoption in residential construction requires a fundamental reorientation of startup commercialisation strategy – away from comprehensive platform plays that maximise addressable market and toward narrow, problem-specific solutions that solve one thing perfectly for a specific workflow. This is not a market positioning observation – it is a structural requirement produced by the Asymmetric Adoption Threshold. Solutions that attempt to transform entire workflows fail because their disruption cost exceeds their visible benefit. Solutions that target one specific pain point succeed because the adoption process does not compete with project delivery. One innovator described their commercialisation philosophy as a direct response to the adoption failure mode they observed across the sector: “We started out being a problem-first company... I think a key failure for a lot of technology companies is they develop something. It might be better, faster, cheaper, but it's looking for a problem to solve... if you don't develop this technology with the construction industry in mind, like from day one, it's probably going to fail.” (I3, p. 272)**:** The problem-first orientation is not just a commercial strategy – it is the mechanism by which the Credibility Threshold is cleared. A solution built from inside a specific workflow demonstrates domain understanding before the first sales conversation begins. This was independently confirmed by a dual-perspective participant who expressed a clear preference for targeted solutions over ambitious overhauls: 'A tangible solution that's an applicable real tool that is solving a single problem and not trying to tackle... every single process of your operation. Simple is better in that way… Too ambitious.' (D1, p. 137) The simpler the solution, the lower the disruption cost and the more clearly the benefit exceeds it. Together, these observations confirm that the path to viable AI adoption in residential construction runs through specificity rather than comprehensiveness – a finding that contradicts the platform logic that dominates startup commercialisation strategy, but that follows inevitably from the structural conditions documented in Themes One through Three.

4.4.3 The Collaboration Imperative

In the residential construction industry, a startup crossing the valley-of-death and reaching successful commercialisation is not only based on product viability, but also on finding the right first-mover. Superior products are created in this space through iterative co-creation on both sides of the interface. Successful ConTech startups often are firms that have found the rare and open-minded incumbent who is willing to co-create, and have embedded themselves in the incumbent's workflows to produce a genuine domain fit. This co-creation is performing the structural function of the missing integrator that the industry fails to do internally. One startup founder described their first successful customer relationship in terms that illustrate the collaborative model precisely: "This guy that I've known for quite a while… (said) I can see the value for my business, why don't you locate (your product) in my facility? He knew it might take a couple of years. He knew he'd have to give up the space for us for free... But now we're actually fabricating products and freeing up his human labour." (I3, p. 280) This is the co-creation model producing the desired outcome for both parties – freeing up human labour for the firm, and adoption for the startup. The incumbent's willingness to absorb short-term cost for long-term benefit is the defining characteristic of the open-minded partner. This is not risk aversion to overcome – it is a different institutional logic operating within the same industry. A second customer encounter produced an equally striking illustration: "This guy is old school. I don't even think he has a high school education, you know, 30 years as a framer. He's like, this is the future. What can I do to help? And it's like, it just, it's unbelievable." (I3, p. 281) The analytical significance here is that domain expertise on the incumbent side produced openness to innovation rather than resistance to it. The innovation is welcomed by a thirty-year framer because the value is inherent to them more than it could be to someone without domain expertise. This is the Credibility Threshold operating in reverse – the incumbent's domain expertise makes them a more valuable co-creation partner, not a more resistant adoption target. A dual-perspective participant explained that finding these open-minded incumbents is difficult because seldom do firms outright say they are against innovation, but that "learning about their (incumbent firms) operations and like their plans and vision." (D1, p. 133) can help differentiate between genuine interest in innovation and those practising tokenism.

Once an open-minded incumbent is identified, the interaction-driven model produces a better domain fit through co-creation – a process that is distinct from Software as a Service (Saas) practices. One innovator described their early adoption strategy as deliberately non-disruptive: "You guys keep your own workflow, but use us alongside of your workflow... that early proof of concept... helped us internally exactly what we need to build out, what we need to fix, change, edit, update."(I1, p. 12) This co-creation process is what produces a domain fit the Credibility Threshold demands. The product improves through use because the incumbent's operational reality is molding its development in real time. The same innovator explained that fundamental to their process was an internal shift from pitching to asking: “It was like, you know, kind of what problem can we solve for you guys as opposed to, you know, just pitching them a solution.” (I1, p. 14) This shift from pitching to asking is the commercial manifestation of the integrating function. In this case, the innovator is performing the structural role that the industry cannot generate internally by embedding itself within the incumbent's operations.

Collaboration is not just beneficial; it is a response to the institutional void identified in Chapter 2. The industry is missing a dedicated systems integrator, and so that role is being filled – in limited and fragile ways – by a combination of open-minded incumbents and synergistic innovators. One dual-perspective participant described their intention to transition into a systems integration role explicitly: “I'm going to become that systems integrator... I have a deep understanding of construction and real estate development... I know how to hire those people.” (D2, p. 88) This is the missing integrator role identified in Chapter 2, not as a theoretical prescription but as an emergent market response. A dual-perspective actor with deep domain expertise on both sides of the interface is independently positioning themselves to perform the structural function that the industry cannot generate internally. The participant's framing is precise – the integrating role requires not technical omniscience but domain-grounded systems coordination: knowing enough construction to build the right systems and enough about talent to hire what they do not know. This is the Credibility Threshold and the integrating function converging in a single actor – the empirical confirmation that the structural void, while real, is not permanently irresolvable.

What these accounts collectively reveal is that viable AI adoption at the interface is not primarily a technology problem – it is a relational and structural one. The open-minded incumbent exists. The co-creation model works. The integrating actor is beginning to emerge. But these conditions are rare, fragile, and dependent on individual relationships rather than systemic structures. The framework developed in Chapter 5 is designed to make these conditions reproducible rather than exceptional.

4.4.4 The Future of Work Divergence

Across the interviews, the topic of the future of work and AI’s role in that transformation was discussed frequently. However, the responses varied widely on how, when, and where AI technology would change residential construction. The industry has not yet developed a shared mental model of what AI transformation looks like – which is itself a barrier to adoption. AI technology is a moving target, both technologically and culturally. Without a consistent imagination across the industry, aiming, investing, and planning for AI is a challenge. An innovator took the conservative position, saying, "Yeah, I would not expect big changes in the next two to three years. It's a slow-moving industry, always has been. Like a lot of industries, it's kind of 1 funeral at a time. So yeah, I'm not like super optimistic about like really, really fast changes" (I4, p. 294). The timeline of this transformation is a furious debate amongst interviewees, with one dual participant predicting a combination of robotics and computer vision would " take over the backbreaking jobs… in 5-10 years." (D1, p. 137) Whereas another dual-perspective participant dismissed robotics entirely, saying, "I think we're 2 decades away from robotics being in and construction. I don't think this whole like talk about robotics and construction. I think it's bull-crap. Like how do robots deal with like weather, right?" (D2, p. 71)The inability of industry members to agree on a common timeline is significant, because without this common understanding, the decisions on when to adopt and when not to will be inconsistent across the industry.

One incumbent's prediction extended beyond incremental change to a fundamental restructuring of labour and ownership: "You're going to be able to utilise AI to make money for you... We're going to have communism through capitalism. So capitalism is going to create communism" (P1, p. 40) Whether or not this prediction proves accurate, its articulation by an active industry practitioner reveals the cognitive range within which participants are imagining AI's ultimate impact – a range so wide that collective institutional preparation becomes structurally difficult to coordinate. This prediction was in stark contrast to other incumbent participants, like one who said: "That's my fear about too much AI, is you're going to lose the human control... and you're going to lose the human contact with the retailer and or with the end user, the homeowner". (P2, p. 178) Whereas another incumbent could not conceive a reality where human workers were replaced on the jobsite: "I think the technologies will help in the background more than the foreground: mapping data, like you said, storing data". But I don't, I don't think it really, I don't see. How AI could make construction better on-site. Really? ... I don't know how you replace a concrete worker or a rebar tyre or steel iron worker" (P5, p. 262). The range of these responses show skepticism and optimism, but more importantly show how fractured the consensus is within the industry.

When skills required for the future of work were discussed one dual-perspective participant explained that collaboration between humans and AI tools is imperative for the next generation of residential construction, saying “Professional prompting is now, now needs to be a job requirement" (D2, p. 109) That professional prompting emerged as a universal short-term skill requirement is significant due to the participants’ skepticism toward AI’s impact. This suggests that human-AI collaboration is the one dimension of transformation the industry can begin preparing for immediately, without resolving the deeper questions of timeline and scale.

Several participants independently identified government-funded pilot programmes as the most viable mechanism for breaking the no-first-mover coordination failure. The innovator with housing policy experience explained a potential method for equipping incumbent firms with the skills necessary to innovate: “There could be a programme or some kind of incentive for these startups and construction companies to come together... maybe they apply jointly and get a certain amount of money to actually do the work together" (I5, pp. 313-314). Without an external actor willing to socialise first-mover risk, the conditional optimism documented throughout Theme Four is likely to remain dependent on rare individual relationships rather than systemic structures.

Taken together, these four themes paint a picture of an industry whose stagnation is not accidental but architecturally produced – sustained by structural incapacities that even willing actors cannot overcome, and by institutional logic clashes that make the commercialisation interface hostile before the first conversation begins. The domain-expert innovator navigating this interface is not merely selling a product but performing a structural function that the industry cannot generate internally. It is this empirical reality that Chapter 5's conceptual framework is designed to address.

5 Conceptual Framework

The empirical findings of Chapter 4 reveal that AI commercialisation failure in Canadian residential construction is not primarily a technology or attitudinal problem but a structural and relational one. A framework response must therefore address the conditions of encounter rather than the characteristics of the technology. *The Collaborative Commercialisation Framework (CCF)* is designed to provide actionable guidance for three audiences simultaneously – innovators navigating the commercialisation interface, incumbents approaching innovation engagement, and policymakers designing the enabling conditions that make productive interface encounters systemically possible rather than accidentally rare. This chapter follows the structure of the framework, introducing each of the four layers sequentially. Section 5.1 presents two diagrams – the first documenting the current hostile interface diagnostically, the second presenting the framework prescriptively – before Sections 5.2 through 5.5 unpack each layer in detail.

5.1 Framework Overview and Design Logic

Standard frameworks treat innovation adoption as a binary outcome – a firm either adopts a technology or it does not. The empirical findings of Chapter 4 demonstrate that this framing is analytically insufficient. Adoption in Canadian residential construction is not a decision but a process. This process is one that requires specific conditions to be met sequentially by multiple actors simultaneously. The CCF is therefore designed not as a decision tool but as a conditions architecture – specifying what conditions must exist at each level of the interface for the level below to become functional. Each layer of the framework's structure reflects this empirical finding directly. Structural conditions must exist before pre-interface requirements can be met. Pre-interface requirements must be met before the interface encounter can be productive. The interface encounter must succeed before institutionalisation becomes possible. The layers do not act as stages that every innovation must pass through sequentially, but are conditions that must simultaneously exist at each level for the level below to function.

Before presenting the framework, it is necessary to establish what the framework is designed to replace. Figure 1 presents the current hostile interface as a system of mutually reinforcing failure conditions – not a collection of independent barriers but a structurally coherent system whose internal logic was never designed to accommodate externally-driven AI commercialisation.

Figure 1: The Hostile Interface: Structural and Institutional Conditions Producing Commercialisation Failure in Canadian Residential Construction

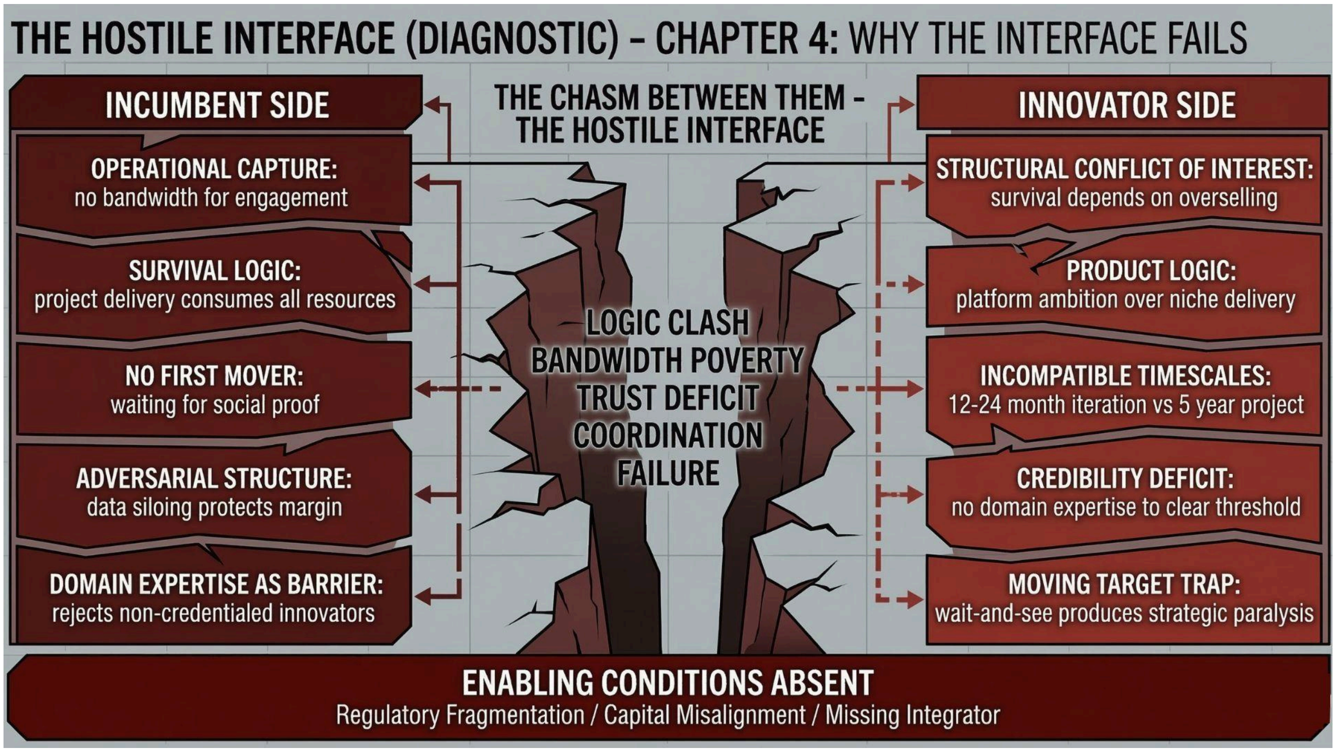

Each failure condition documented in Figure 1 has a direct response in Figure 2. Each element of the framework was designed in response to an empirically documented barrier rather than derived from theoretical prescription alone.

Figure 2 presents the *Collaborative Commercialisation Framework* – a four-layer conditions architecture designed to transform the hostile interface documented in Figure 1 into a collaborative ecosystem within which AI commercialisation becomes viable. The framework addresses three audiences simultaneously – innovators navigating the commercialisation interface, incumbents approaching innovation engagement, and policymakers designing the enabling conditions that make productive encounters systemically possible. Each layer is developed in full in the sections that follow.

Figure 2: The Collaborative Commercialisation Framework: Enabling Conditions and Interface Requirements for Viable AI Adoption in Canadian Residential Construction.

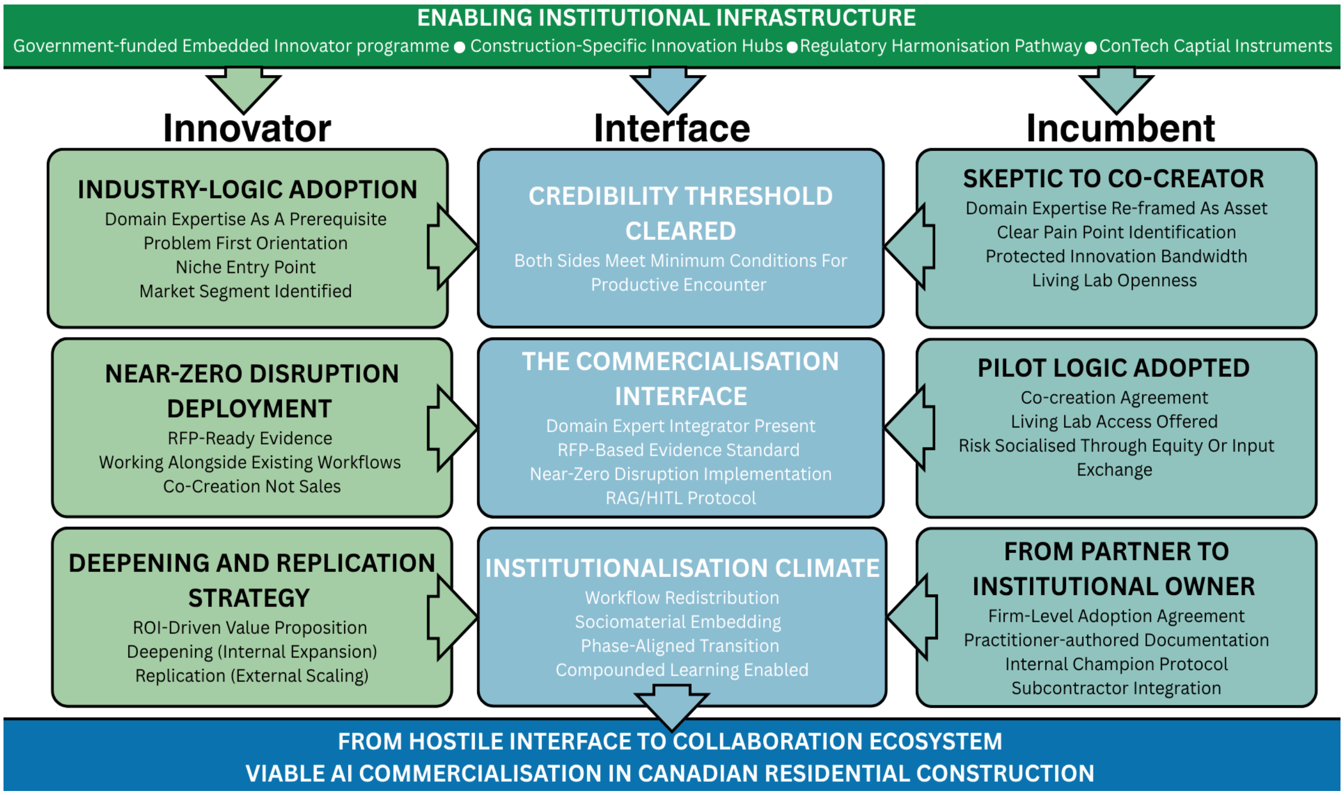

The enabling institutional infrastructure layer is developed in Section 5.2, followed by the pre-interface requirements in Section 5.3, the front-line encounter conditions in Section 5.4, and the institutionalisation climate in Section 5.5.

5.2 Enabling Institutional Infrastructure

The empirical findings in Chapter 4 were clear: behavioural change alone is insufficient when the structural conditions within which the interface exists actively prevent a productive encounter. The enabling institutional infrastructure layer represents the key to change. If changes are made at this foundational layer, it will empower those existing at the interface to take the actions required for AI adoption in Canada's residential construction industry. This section specifies the policy and ecosystem conditions that must exist for the lower layers of the framework to function systematically rather than accidentally. The ecosystem coordination failure documented in Section 4.3.4 established that no individual actor has sufficient incentive to address structural conditions unilaterally – policy intervention is therefore not supplementary to the framework but the precondition for it. The coordination failure ensures that no individual actor has sufficient incentive to address structural conditions unilaterally.

Policy intervention is therefore not supplementary to the framework – it is the precondition for it.

5.2.1 The Embedded Innovator Partnership Programme

The framework recommends that the government fund innovators with demonstrated domain expertise to embed within incumbent firms as systems integrators – enabling them to learn construction operations, iterate their product in real project conditions, and perform the structural integrating function that the industry cannot generate internally. Participants independently identified government-funded pilot programmes as a promising mechanism – D1 noted that CMHC had begun funding pilot projects – though the empirical findings suggest these efforts have not yet produced the systematic interface conditions the framework requires.

This recommendation is motivated by three converging findings – the missing integrator vacancy identified in Chapter 2 and empirically confirmed in Section 4.4.3, the Operational Capture finding of Section 4.3.1, which established that incumbents cannot engage with innovation without bandwidth relief, and the no-first-mover dynamic of Section 4.2.2, which established that the coordination failure requires an external actor to socialise first-mover risk. Independent policymaker confirmation of this specific mechanism was documented in Section 4.4.4 – a former municipal housing executive independently identified an innovator-incumbent pairing programme with equity participation as the primary mechanism for reducing interface hostility. Together, these converging findings show the embedded innovator programme not as a speculative policy proposal but as an empirically motivated response to three simultaneous structural failures: missing integration capacity, incumbent bandwidth poverty, and ecosystem coordination failure.

5.2.2 Construction-Specific Innovation Infrastructure

Canadian innovation hubs and accelerator programmes should develop construction-specific streams with appropriate physical infrastructure – warehouse space, fabrication facilities, materials testing environments, and access to active construction sites – rather than assuming that SaaS-oriented shared office infrastructure meets ConTech startup needs. For AI tools specifically, physical access to construction environments is not a convenience but a data necessity – domain-specific training data that makes AI reliable enough to clear the

Asymmetric Adoption Threshold can only be generated through real project exposure, not simulated in a software-oriented co-working space.

This recommendation responds to the incumbent observation documented in Section 4.3.4 that the current startup infrastructure is structurally misaligned with ConTech's physical requirements. An innovator building AI-powered robotics or materials solutions cannot iterate on their product in a co-working space designed for software companies. Construction-specific innovation infrastructure is the physical prerequisite for the embedded innovator programme – without appropriate facilities, innovators cannot iterate their products in real conditions, regardless of how well the partnership programme is funded.

5.2.3 Regulatory Harmonisation Pathway

Policymakers should establish a national certification pathway allowing AI solutions validated in one Canadian province to be recognised in others without full redesign or recertification – or, more ambitiously, pursue harmonisation of building requirements across jurisdictions. Either approach should preserve provincial jurisdiction while creating mutual recognition frameworks that reduce the regulatory burden of cross-jurisdictional expansion; this pathway need not eliminate provincial jurisdiction over building codes – it should create a mutual recognition framework that allows demonstrated performance in one jurisdiction to reduce the regulatory burden of entry into others.

This recommendation responds to the Regulatory Non-Replicability concept introduced in Chapter 2 and empirically confirmed across multiple participants in Section 4.3.3 – innovators described being geographically locked to their initial jurisdiction, with expansion requiring them to go back to the drawing board entirely. Likewise, incumbents described the labour-intensive process of adapting to building codes within the same province. (Section 4.3.3) The policymaker participant independently identified municipal alignment as the single biggest barrier to scaling modern construction methods in Canada. The National Building Code framework indicates that cross-provincial regulatory coordination is institutionally possible even within Canada's federated structure. Without regulatory harmonisation, innovators who successfully navigate the interface and validate their solution in one jurisdiction remain geographically locked – their commercial success producing a beachhead rather than a scalable business.

5.2.4 ConTech-Specific Capital Instruments

Government and institutional investors should develop ConTech-specific capital instruments with return horizons and milestone structures aligned with construction's adoption pace – patient capital funds, government co-investment programmes, and innovation tax credits for construction firms investing in digital transformation – rather than applying VC return logic that is structurally incompatible with ConTech commercialisation timescales.

VC firms follow herd behaviour rather than sector need, and their return requirements are structurally misaligned with the niche-first scaling strategy that the empirical findings support. Innovators who follow the problem-first, near-zero disruption approach documented in Section 4.4.2 cannot simultaneously deliver the rapid market expansion VC firms require. Without alternative capital instruments, the commercialisation strategy the framework proposes remains commercially unviable regardless of its technical and relational merits.

5.2.5 Regulatory Agility for New Materials and Methods

Regulatory bodies should develop adaptive assessment frameworks for novel construction materials and methods. This should include AI-optimised prefabrication and modular construction, so that new approaches can be evaluated on performance outcomes rather than prescriptive compliance with codes designed for conventional construction methods.

This recommendation responds to empirical findings in Section 4.3.3 – participants identified that regulators lack the capacity or methodology to fold new technologies into existing approval frameworks, and that the absence of a regulatory framework supporting industrialised construction represents a systemic barrier to the productivity transformation AI could enable. Without regulatory agility at the approval level, AI solutions that optimise prefabrication and modular construction methods cannot be deployed at scale, regardless of their demonstrated performance – the regulatory framework itself becomes the ceiling on what the industry's digital transformation can achieve.

These five enabling conditions are interdependent rather than independent policy prescriptions. The embedded innovator programme requires construction-specific physical infrastructure to function because innovators cannot iterate in real conditions without appropriate facilities. Regulatory harmonisation makes the validated solution scalable beyond the initial jurisdiction. ConTech capital instruments give the innovator the runway to complete the niche-first commercialisation strategy without abandoning it for VC-compatible growth logic. And regulatory agility for new materials removes the ceiling on what AI-enabled construction methods can ultimately achieve. Together, they constitute the institutional preconditions for the framework's remaining layers to function systematically and transform the conditions that make productive interface encounters rare into the conditions that make them reproducible.

5.3 Pre-Interface Requirements

The enabling institutional infrastructure is the key facilitating factor for creating a more frictionless interface. Once the institutional infrastructure is no longer Architecturally Hostile, the path is cleared for incumbents and innovators to approach the interaction interface. The removal of structural hostility is a necessary but insufficient condition for productive interface encounters. Both actors must arrive at the interface having met specific pre-conditions that determine whether the Credibility Threshold can be cleared and co-creation can begin. This framework outlines the specific preconditions that must be met on either side of the interface to produce a positive interaction. These are not personality requirements or attitudinal preferences, as much as they are structural prerequisites. The empirical findings of Chapter 4 demonstrate that when productive encounters do occur under current conditions, they are rare, fragile, and dependent on individual relationships rather than systemic structures. The pre-interface requirements specified here are designed to make productive encounters reproducible rather than exceptional.

5.3.1 The Innovator Side: Industry Logic Adoption:

At this stage, it is imperative for the innovator to approach the interface not as an AI technology company selling a product but as a domain-fluent problem solver. The innovator must possess construction domain expertise – either through their own background or through deliberate team composition – as the Credibility Threshold finding established that this is a prerequisite condition for interface engagement rather than a competitive advantage. The innovator should speak project logic, not product logic, presenting themselves as a partner solving one specific operational pain point with a clear impact on saving time or money, rather than a systematic technology set to transform the industry. At this stage, the innovator should act as a subcontractor would – borrowing the institutionalised trust structure that incumbent firms already understand – ready to respond to scope rather than pitch, and prepared to answer an RFP rather than deliver a demo. This shift of language and commercial behaviour clears the Credibility Threshold by design rather than by accident.

Before approaching the interface, however, the innovator must make a deliberate strategic decision about which segment of the bifurcated market to target first. The empirical findings of Chapter 4 established that Canadian residential construction is divided into two different adoption failure modes: micro-firms whose survival logic and bandwidth poverty foreclose meaningful engagement, and large firms whose analysis paralysis and organisational inertia prevent commitment. Attempting to address both segments simultaneously is a commercialisation strategy that the structural conditions documented in Chapter 4 make untenable. The framework recommends that innovators identify and target the small and mid-size firm segments as their primary beachhead. These are firms large enough to possess some innovation capacity and protected bandwidth, but small enough to make decisions without the bureaucratic inertia that afflicts larger organisations. This segment is not immune to the barriers documented throughout this thesis, but it represents the viable entry point where the *Embedded Value Escalation Strategy* is most likely to clear the Asymmetric Adoption Threshold. Once a reference case is established within this segment, the evidence produced through co-creation provides the credibility and performance metrics required to approach both micro-firm and large-firm segments with dramatically reduced friction.

5.3.2 The Scaling Paradox

It must be addressed that the niche-first approach on the side of innovators conflicts with VC growth logic. VC capital will be difficult to obtain for innovators applying niche solutions and running pilot projects. ConTech's capital instruments present a problem that must be solved through structural response by top-down policy adjustments. Alternative capital instruments must be available that reward niche-first scaling rather than penalising it. This is a structural barrier presented in 4.3.4 that cannot be easily or reliably reconciled by those at the interface. The deepening and replication pathways described above partially mitigate this tension, as validated reference cases and documented performance metrics make ConTech investment marginally more attractive than an unproven niche product. But this mitigation is insufficient without structural capital reform.

5.3.3 The Incumbent Side: Sceptic to Co-Creator

Step One – Domain Expertise Reframed

Equally as important in a positive interface interaction is a reframe of the incumbent from a sceptic to a co-creator. The incumbent's domain expertise in construction is not a standard against which innovations are judged and found wanting, but instead, it is a resource used to inform innovation. This is not merely an attitudinal recommendation, but it is a structural one. The Credibility Threshold finding established that domain expertise determines whether the interface functions at all. But the co-creation evidence from Section 4.4.3 extends this finding – incumbent domain expertise does not just clear the threshold, it actively produces the domain fit that makes AI tools viable in construction conditions. An incumbent who withholds their operational knowledge as an evaluative standard is structurally preventing the co-creation process that would make the innovation worth evaluating. Innovators without access to domain expertise build solutions that fail at the interface because they fail the sim-to-real test. The framework recommends that incumbents approach the interface as active contributors rather than passive evaluators. This reframe is empirically grounded in the Credibility Threshold operating in reverse – documented in Section 4.2.1, where a thirty-year framer's domain expertise produced openness to innovation rather than resistance, because he understood the value the innovation offered in terms of his experience made legible.

Step Two – Pain Point Identification and Change Management Budget.

Before engaging with any innovator, the incumbent should be able to clearly communicate workflows or processes they would pay to improve. This is both the filter for evaluating innovations and the brief that guides productive co-creation. This recommendation is grounded in the niche solution finding of Section 4.4.1, where incumbents across the sample were actively searching for targeted solutions to specific operational problems but struggling to articulate those problems accurately enough to guide innovators toward viable solutions. Alongside pain point identification, incumbents should allocate a modest protected budget for innovation engagement – even a minimal designated resource signals organisational commitment and creates the bandwidth that Operational Capture currently forecloses. These two components are interdependent – pain point identification without protected bandwidth produces a brief that nobody has time to act on, and bandwidth without a focused brief produces unfocused experimentation that Operational Capture will eventually reclaim.

Together, they constitute the minimum organisational conditions for productive interface engagement.

Step Three – Living Lab Partnership

The framework recommends that incumbents enter the interface with a prior organisational commitment to treating their construction sites as living labs – a component of the broader innovation infrastructure the industry currently lacks. Incumbents arrive at the interface prepared to offer site access and operational knowledge in exchange for equity, revenue sharing, preferential rates, or co-development credit. This aligns incentives structurally – the incumbent has a genuine financial stake in the innovation's success and therefore will be more willing to co-create with the innovator. The exchange socialises first-mover risk through compensation rather than requiring altruism, and it addresses the missing integrator vacancy by embedding the innovator within the construction environment where domain fit is produced. The specific terms of the living lab arrangement – project phases, equity structure, co-development credit, and exit conditions – are negotiated during the interface encounter and form a core component of the co-creation agreement developed in Section 5.4. What this pre-interface step requires is not the agreement itself but the organisational disposition that makes the agreement possible.

5.4 The Front Line: Trenches to Treaties

Once the pre-interface requirements are met by both the innovator and incumbent, the encounter is still at risk to fail. The logic by which the interaction is dictated can be the catalyst that determines whether AI technology bridges the commercialisation gap. Construction's default procurement logic enforces competition – lowest bid, strategic information withholding, transactional relationships. AI commercialisation requires interaction in the form of collaborative co-creation, transparent data sharing, and sustained relational trust. The Front Line layer specifies the conditions under which the interaction logic can replace the competition logic at the interface.

5.4.1 Innovator Side: Stage Two – Near-Zero Disruption Implementation:

Once both sides have met their pre-interface requirements and entered the encounter, the innovator's primary obligation is to deploy their solution in a way that demonstrates value without threatening project delivery – the Near-Zero Disruption Implementation stage of the

Embedded Value Escalation Strategy. Innovations should initially target non-disruptive sections of the incumbent's workflow – processes that can be tested in parallel without competing with active project delivery. Once an incumbent firm has decided to adopt the innovator's proposal, the solution is to be deployed alongside existing workflows rather than replacing them. The incumbent maintains current practices while the innovation proves value in parallel. This ensures that innovators have the safety-net of failure on first implementation without burning relations with thin-margined incumbent firms. This allows innovators to learn with real-project data and for incumbents to witness potential value and contribute key insights that can improve product fit. This shifts innovators' role from pitching/promising value to producing, which creates trust and reduces the possibility of overselling innovations. This shift from pitching to producing resolves the structural conflict of interest identified in Section 4.2.4 – the innovator whose value is demonstrated in real conditions has no structural incentive to oversell because the product's performance replaces the sales conversation. Adoption cost is engineered to near zero, so the Asymmetric Adoption Threshold becomes clearable. The co-creation process begins – the incumbent's operational reality shapes the product's development in real time. This approach is empirically grounded in the co-creation model documented in Section 4.4.3, where innovators who deployed alongside existing workflows rather than replacing them produced genuine domain fit while building the trust that subsequent adoption required.

5.4.2 Incumbent side: The Living Lab

The living lab is the place where the incumbent transforms from the passive evaluator of AI tools to the active co-creator of innovation. This arrangement replaces the traditional sales agreement that is centred around a finished product with a co-development agreement that binds both parties into building a valuable innovation while protecting project delivery. Project delivery must be protected for the incumbent; the innovator risks disrupting timelines and ROI for the incumbent firm, which could prove catastrophic for the interface.

The living lab arrangement should be structured first around low-disruption project phases – design completion, pre-construction, and post-occupancy, rather than active delivery phases where failure costs are highest. The parallel deployment strategy established in Stage Two of the Embedded Value Escalation Strategy provides the structural mechanism for this protection – the innovation operates alongside existing workflows during active delivery phases, transitioning to primary deployment only at phase boundaries where disruption cost is lowest. The living lab and the parallel deployment strategy are therefore complementary rather than independent recommendations – each depends on the other for its viability.

Three viable options for a compensation structure between innovators and incumbents exist, and the agreement should specify which applies. Equity participation – the incumbent receives a small ownership stake in the startup in exchange for site access and operational knowledge contribution. Revenue sharing – the incumbent receives a percentage of revenue generated from solutions developed using their site data and operational input. Co-development credit – the incumbent receives preferential pricing, first-access rights to new features, or named co-development acknowledgement. Connect to I5's independent identification of equity as the mechanism for securing genuine organisational commitment rather than performative engagement. Beyond compensation, the agreement must address intellectual property provisions explicitly. The operational knowledge and site data that the incumbent contributes to co-creation have commercial value. The agreement must specify how this contribution is recognised – who owns the data generated during the living lab period, how domain-specific configurations developed during co-creation are attributed, and what restrictions apply to the innovator's use of incumbent-specific operational knowledge in products deployed to competitor firms. This is a protection for the incumbent that makes the living lab arrangement commercially safe rather than commercially naive.

Finally, both parties need clearly defined conditions under which the living lab arrangement can be terminated without damaging the underlying business relationship. If the innovation fails to demonstrate value within a specified phase, the incumbent can exit without penalty. If the incumbent's operational conditions change – new project pressures, personnel turnover, financial stress – the arrangement can be suspended rather than abandoned. Clean exit conditions paradoxically make commitment easier because they remove the fear of being locked into a failing experiment.

The living lab agreement transforms the interface encounter from a transactional sales event into a governed co-development relationship. Its four components address the four specific failure modes documented in Chapter 4 – Operational Capture through phase-specific access, first-mover risk through compensation structure, data exploitation risk through IP provisions, and commitment anxiety through exit conditions. Together, they constitute the structural mechanism through which the Competition Versus Interaction Paradox is resolved at the individual firm level.

5.4.3 The Domain Expert Integrator

The interface needs an integrator who can translate between the construction project's logic and the startup's product logic. This integrator must be able to calibrate the co-creation process against site conditions and perform the structural function that the industry cannot generate internally. This is the missing integrator role identified in Chapter 2, and then empirically confirmed as beginning to emerge organically in Section 4.4.3. The integrator can emerge in the following three forms:

The domain-expert founder: The innovator themselves possesses sufficient construction domain expertise to perform the integrating function. This is the Credibility Threshold cleared at the individual level – the Grey Hair Effect operating through the founder. This is the most powerful form because the integrating function is structurally embedded in the commercialisation relationship from the start.

The embedded team member: The innovator's team includes a construction domain expert whose role is specifically to perform the integrating function during co-creation. This form is empirically confirmed by P1's decision to hire an experienced PM specifically to bridge the credibility gap with clients – a move that immediately transformed the firm's ability to engage at the interface.

The external systems integrator: A dual-perspective actor operating independently – with deep construction domain expertise and sufficient technical literacy to bridge both logics – performs the integrating function as a third-party service. This emergent form is empirically confirmed in Section 4.4.3, where a dual-perspective participant with deep construction and development expertise independently identified the systems integrator role as both a market opportunity and a structural necessity. Hoping that they could position themselves to perform the function that the industry cannot generate internally. This is the form that government-funded embedded innovator programmes are best positioned to support – connecting domain-expert integrators with startups that lack construction backgrounds and incumbents that lack digital literacy.

Without this actor present, the bandwidth-constrained incumbent has no viable pathway to informed adoption, and the Credibility Threshold remains uncleared regardless of the innovation's technical merit. The domain expert integrator is not just a helpful addition to the commercialisation encounter, but they are the structural mechanism through which the

Credibility Threshold is cleared, the cognitive offloading function is performed, and the institutional logic translation that makes co-creation possible is executed.

5.4.4 Co-Creation Agreement Structure

The living lab agreement governs the physical and commercial terms of the co-development relationship. The co-creation agreement governs its intellectual and operational terms. This agreement specifies how both parties will work together, what each will contribute, how decisions will be made, and how the outputs of co-creation will be owned and used.

The innovator contributes technical development capacity, AI tool iteration, and domain-specific product configuration. The incumbent contributes site access, operational knowledge, project data, and practitioner feedback. Both contributions should be documented explicitly – not as a legal formality but as a commitment mechanism that makes the co-creation relationship legible to both parties. Product updates and new feature deployments should be timed to project phase boundaries rather than the innovator's internal sprint cycles. This prevents the innovation from disrupting active delivery while ensuring that the incumbent's operational feedback reaches the product development process in time to influence the next iteration. This agreement would also indicate who is responsible for deciding when the co-creation process has succeeded or failed. Clear decision rights prevent the relationship from stalling at the transition between Stage Two and Stage Three – the moment where the evidence base from parallel deployment should trigger controlled full deployment. The co-creation agreement should specify that operational documentation – written by practitioners, capturing what works and what does not in construction-specific conditions – is a deliverable of the co-creation process rather than an optional output. Without this specification, the documentation that makes institutionalisation possible is never produced. Co-creation relationships generate ambiguity about responsibilities and what constitutes success. A lightweight dispute resolution mechanism – not legal arbitration but a named third-party mediator, ideally the domain expert integrator – prevents commercial disagreements from destroying the collaborative relationship that makes the innovation viable.

Together, these components transform the co-creation agreement from a legal necessity into an operational architecture – the shared agreement that makes sustained collaboration possible across the project-based discontinuity that would otherwise dissolve the relationship when the project closes.

5.4.5 Preventive Measures for AI Reliability at the Interface

The hallucination trust deficit documented in Section 4.2.3 cannot be resolved before the interface encounter because it requires domain-specific data that only co-creation can generate. The Front Line layer, therefore, includes two complementary technical conditions for managing hallucination risk during co-creation: the Retrieval Augmented Generation (RAG) architecture and the human-in-the-loop protocol. AI outputs are informed by hand-picked data sets and reviewed by a domain-expert human before being acted upon. These measures do not eliminate the hallucination problem, but they calibrate incumbent risk tolerance to a manageable level.

During the co-creation process, the innovator and incumbent together define the closed-loop data pool that constrains the AI's outputs. Rather than drawing from open internet sources, the AI is restricted to verified domain-specific information – zoning data, building codes, project specifications, historical cost data – that both parties have validated as reliable. This approach is empirically grounded in Section 4.2.3, where I5 noted that hallucinations were essentially eliminated by restricting the AI to pre-verified domain data, and D1 independently observed that numerical and rules-based inputs are less prone to hallucination than interpretive ones. This suggests that the initial RAG data pool should prioritise structured data before incorporating unstructured operational knowledge. Notably, the setup of this system is solution-specific and will vary in terms of bandwidth requirements. The RAG architecture is a co-creation output rather than a pre-interface preparation because the relevant domain data only becomes available through the incumbent's operational knowledge contribution during the living lab process. This is why it belongs in the Front Line rather than the pre-interface layer.

The second complementary condition is the human-in-the-loop protocol. All AI outputs during the co-creation phase should be reviewed by a domain-expert human before being acted upon. It manages hallucination risk at the operational level, and it generates the practitioner feedback that improves the RAG data pool over time – the human reviewer who catches an error is simultaneously identifying a gap in the domain-specific data the AI has been given. That feedback loop is the mechanism through which the hallucination trust deficit is progressively resolved through co-creation rather than deferred to future technological improvement.

Together, the RAG architecture and human-in-the-loop protocol transform the hallucination trust deficit from a structural barrier into a value-added process. Rather than waiting for AI reliability to improve to the point where incumbent risk tolerance is satisfied – the Moving Target Trap dynamic identified in Section 4.2.3 – these conditions make reliability improvement an output of the co-creation relationship itself. The stability being waited for is produced through the encounter rather than arriving from outside it.

The conditions specified in this section collectively constitute the interaction logic that replaces construction's default competition logic at the commercialisation interface. No single condition is sufficient alone. The living lab agreement without the domain expert integrator produces a co-development relationship that neither party can navigate. The RAG architecture without the human-in-the-loop protocol produces a reliability mitigation without the feedback loop that makes it improve. Together, these conditions produce the boundary-spanning encounter that transforms the hostile interface documented in Chapter 4 into the collaborative ecosystem the framework is designed to enable.

5.5 Institutionalisation Climate

5.5.1 Incumbent Side Stage Three: Controlled Deployment and Institutionalisation

The final stage is where the innovation stops being tested alongside existing practices and becomes the primary system. The timing of this transition is critical – it should happen at a project phase transition point rather than mid-delivery, and it should be scoped to one function at a time rather than attempted as a wholesale workflow replacement. The innovator uses performance metrics gathered during parallel deployment to demonstrate exceeding value over the cost of disruption. Once the value proposition is clearly laid out as an ROI opportunity, the full-deployment can be absorbed because the evidence base exists.

To ensure the innovation adoption survives the project-based continuity and one-off project teams, the innovator documents knowledge in conjunction with the practitioners. Domain-specific protocols and workarounds must be documented in a useful way that survives personnel turnover. This documentation imperative takes three specific forms:

operational documentation, internal champion identification and firm-level adoption agreement. Operational documentation is not technical documentation written by the innovator, but written by the practitioners who used it, capturing what it does well, what it does not do well, and how it integrates with construction-specific site conditions. An internal champion is identified, one or two people who genuinely understand the tool and advocate for it. The framework recommends deliberately identifying these champions early and giving them formal responsibility for knowledge transfer to the next project team. Without a named champion, the institutional knowledge dissolves when the project closes. The co-creation agreement from Stage Two governs a specific project. Stage Three requires a firm-level adoption agreement that commits the incumbent organisation to continued use. This shifts the adoption from a project experiment to an organisational practice. This recommendation responds to the project-based dissolution dynamic identified in Section 4.3.2 – the same structural condition that prevents data from carrying forward between projects threatens to dissolve successful innovation pilots when project teams disperse. Once the innovation is embedded at the firm level through these three mechanisms, the trust and operational familiarity established through Stages One and Two make broader digital transformation engagement possible – not as a sales proposition but as a natural extension of a proven relationship.

Successful Stage Three institutionalisation produces two viable commercial pathways for the innovator. The first is deepening – expanding the innovation's footprint within the same firm through adjacent solutions and broader digital transformation engagement, made possible by the trust and operational familiarity the three-stage process has produced. The second is replication – bringing the co-created, field-validated product to new firms with a reference case, documented performance metrics, and practitioner-authored operational protocols that dramatically reduce the Credibility Threshold barrier in subsequent commercialisation conversations. Together, these pathways reframe the niche-first strategy not as a permanent commercial limitation but as a sequenced market development approach – one that trades short-term addressable market size for the domain fit, trust, and validated evidence base that makes sustainable scaling possible.

5.5.2 Subcontractor Integration Protocol

The framework recommends that firms make subcontractor participation with adopted AI tools a contractual obligation within their project agreements. A firm-level adoption agreement means nothing if the subcontractors whose activities generate the most operationally critical data refuse to engage with the digital tool. Institutionalisation at the firm-wide project level without subcontractor integration produces a system that works in the office and fails on-site. The framework recommends that co-creation agreements include a subcontractor engagement plan – identifying which subcontractors must use the tool for it to deliver value, how their participation will be incentivised, and what minimum digital literacy requirements apply. This recommendation responds to the subcontractor dependency finding of Section 4.3.2 – even willing incumbents cannot build the data infrastructure AI requires when subcontractors strategically silo information to protect their bargaining position. Without this protocol, the institutionalisation mechanisms specified in Stage Three remain vulnerable to the same subcontractor fragmentation that prevents data infrastructure development across the industry.

5.5.3 From Boundary Object to Embedded Practice

The ultimate measure of successful institutionalisation is not continued use, but the transition from boundary object to embedded practice. A tool that is used because a champion advocates for it is still a boundary object, and its status is dependent on individual relationships and vulnerable to personnel turnover. A tool that has been embedded in the firm's operational protocols, documented by practitioners, and adopted at the firm level rather than the project level has become an institutional practice – its continuation is no longer dependent on any individual's advocacy. Orlikowski’s (2007) Sociomateriality theory establishes that technology becomes institutionalised when it redistributes authority and workflow in ways that become taken for granted – when the question is no longer whether to use the tool but how to use it better. The documentation, champion identification, and firm-level agreement mechanisms specified in Stage Three are the practical pathways to this sociomaterial embedding. This transition – from hostile interface to co-created pilot, to embedded institutional practice is the full arc the CCF is designed to enable, and the measure against which its success should ultimately be evaluated.

5.6 Framework Limitations and Boundary Conditions

The CCF is grounded in empirical evidence and designed to address documented structural conditions, but its scope, sample, and design logic carry limitations that define the conditions of its applicability.

5.6.1 The simultaneous behaviour change requirement

The framework requires innovators, incumbents, and policymakers to all change simultaneously. In practice, the width and depth of these changes are unlikely to occur without a coordinating actor willing to absorb the transactional costs required to align these independent actors. The embedded innovator programme partially addresses this by providing a coordinating mechanism at the firm level – but ecosystem-level coordination remains dependent on policy commitment that the framework cannot guarantee. This is the framework's most significant implementation challenge.

5.6.2 The scaling paradox

The framework recommends a niche-first strategy for innovators as a prerequisite for initial adoption. This creates a structural tension between the niche-first entry pathway and the rapid market expansion that startups require for investor confidence and long-term commercial viability. The enabling layer partially addresses this by recommending ConTech-specific capital instruments, but cannot fully resolve it through framework design alone. This scaling paradox will ultimately require market-level changes in how ConTech investment is evaluated and funded that extend beyond the scope of any single framework's prescriptive reach.

5.6.3 Sample Size and Context Boundary

The framework is derived from twelve 45-minute interviews conducted with participants operating in Canadian residential construction. Canada's residential construction market is geographically diverse – the regulatory, climatic, and market conditions facing builders in Toronto differ meaningfully from those in Calgary, Vancouver, or rural Nova Scotia. A twelve-interview qualitative study cannot claim national representation, and its value therefore lies in theoretical depth rather than statistical breadth. Its recommendations are empirically grounded propositions for further validation, not proven prescriptions applicable to all construction contexts. The framework's applicability to commercial construction, international markets, or non-AI digital technologies requires empirical testing beyond the scope of this research. The structural and institutional conditions it addresses are not unique to Canada and may have transferable relevance to other fragmented, project-based construction markets facing similar AI commercialisation challenges.

5.6.4 Moving Target Trap

The Moving Target Trap identified in Chapter 2 operates not just at the level of incumbent adoption decisions but at the level of framework design itself. This framework was designed to address AI commercialisation as it exists at the time of this research – characterised by the specific reliability profile, capability limitations, and institutional unfamiliarity documented in the empirical findings. If AI capabilities shift dramatically – if hallucinations become negligible, if AI achieves genuine autonomous site operation, if large language models develop genuine construction domain knowledge without human input – some of the framework's specific recommendations may become obsolete, particularly those addressing hallucination risk and domain expertise requirements. The framework's structural logic is designed to be robust to technological change. Aiming to address conditions of encounter rather than characteristics of technology. However, its specific operational recommendations should be treated as temporally bounded propositions requiring periodic reassessment as the AI landscape progresses.

5.6.5 The micro-firm bandwidth constraint

The framework's pre-interface requirements assume a minimum level of organisational capacity that the smallest micro-firms may not possess. The protected innovation bandwidth recommendation, the pain point identification process, and the living lab partnership model all require slack resources that firms operating at survival logic margins cannot generate. The enabling layer's embedded innovator programme and capital instruments partially address this problem, but micro-firms below a certain operational threshold may remain outside the framework's effective boundary regardless of policy intervention. This limitation is consistent with the bifurcated market finding of Chapter 4 and does not invalidate the framework, but it defines its primary target as mid-size and growth-oriented firms rather than the full residential construction market. Ideally, once the innovation has proven its value, it will trickle-down into smaller firms.

Together, these limitations do not undermine the framework's analytical contribution – they define the conditions under which it applies and the empirical questions that future research must address to extend its reach.

6 Evaluation and Falsification

The CCF proposes a structural response to a documented systemic failure. This chapter evaluates whether that response is analytically justified, internally coherent, and transferable from theoretical design to practical implementation. Methodologically, two complementary methods are applied. The Sim-to-Real transferability lens assesses the gap between the framework's controlled design conditions and the messy structural reality of Canadian residential construction. The second method is a literature-based stress test that applies the theoretical toolkit established in Chapter 2 to identify where the framework's assumptions are most theoretically vulnerable.

Section 6.1 establishes the evaluative approach. Section 6.2 presents the strongest transferability evidence. Section 6.3 applies four structured stress tests. Section 6.4 delivers the net transferability verdict. 6.5 identifies implications for further research.

6.1 The Evaluative Approach

6.1.1 The Sim-to-Real Lens

In engineering and robotics, the Sim-to-Real gap describes the difference between how a system performs under controlled simulation conditions and how it performs in the variable, unpredictable conditions of real-world deployment. Applied to the conceptual framework evaluation, the same logic holds. A framework designed from qualitative interview data represents a controlled theoretical environment, and the real world of Canadian residential construction is messier, more variable, and more structurally resistant than any theoretical model can fully capture. The Sim-to-Real lens asks three structured questions. Does the framework capture the relevant variables – are all documented barriers addressed? Does it make realistic assumptions about actor behaviour – does it require behaviour that the structural conditions documented in Chapter 4 make unlikely? Lastly, does it account for environmental variability – does it hold under different market conditions, firm sizes, and regulatory environments?

6.1.2 The Literature Stress Test

A complementary evaluative method applies the theoretical toolkit from Chapter 2 to identify where the framework's assumptions are most theoretically vulnerable. Four stress tests are constructed – each identifying a specific point where an established theoretical framework would predict the CCF to face transferability challenges. Each stress test is followed by the framework's response and an honest assessment of the residual gap. This method treats the framework not as a theoretical conclusion but as a theoretical proposition that must demonstrate relation with the body of knowledge it draws from before it can claim analytical credibility.

6.2 Transferability Strengths: Where the Framework Holds

The framework's most important transferability strength is that its core mechanisms were not derived from one side of the interface alone. The Credibility Threshold finding, the no-first-mover dynamic, the adversarial structure diagnosis, and the collaboration imperative were all independently confirmed by participants on both sides of the interface. This bilateral confirmation suggests the framework reflects structural conditions of the interface rather than the positionally biased perceptions of any single actor type.

Several of the framework's most original mechanisms were independently identified by participants without theoretical prompting. The systems integrator role was named and actively pursued by D2 without knowledge of the missing integrator concept from Chapter 2. The government pairing programme was independently prescribed by I6 without knowledge of the framework's enabling layer. The equity mechanism for securing genuine incumbent commitment was named by I6 independently. The problem-first commercialisation approach was described by multiple innovators as their actual strategy rather than a theoretical recommendation.

This convergence between theoretical framework design and practitioner independent identification is the strongest possible evidence of real-world transferability – the framework is not prescribing behaviour that practitioners find implausible, it's systematising behaviour that successful practitioners have already discovered through experience. The independent confirmation from I6 – a former municipal housing executive with direct responsibility for housing policy – extends the framework's evidential base beyond commercial actors to include a policymaker perspective. I6's independent identification of the pairing programme and equity mechanism as primary solutions to interface hostility constitutes external validation from a non-commercial actor with direct implementation experience. This confirmation is particularly significant for the enabling layer recommendations and the sections of the framework most vulnerable to the charge of being theoretically motivated rather than practically grounded.

6.3 Stress Tests: Where the Framework Faces Transferability Challenges

6.3.1 Stress Test One: Living Lab vs Survival Logic

Introduced in Chapter 2, Institutional logics (Thornton and Ocasio, 2008) would predict that when project stress intensifies, the dominant project and survival logics that dominate the residential construction industry will reemerge. The reemergence of these logics will come at the expense of the innovation logic that the living lab recommendation requires. The framework relies on the concept of the living lab being resilient enough and providing enough value to firms and innovators to survive the conditions that induce project stress: margin pressure, delivery stress, and Operational Capture. Empirically, these stresses were found to be endemic to the industry, and thus, cannot be structurally avoided by any framework recommendations. The first stress test attempts to uncover whether the framework provides the structural mechanisms to protect the living lab when incumbents revert to survival mode.

Three structural mechanisms within the framework address this specifically. First, the phase-specific access rights to the living lab agreement restrict co-creation activity to low-disruption project phases. The recommended phases – design completion, pre-construction, post-occupancy – are structurally outside the active delivery pressure that triggers Operational Capture and survival logic reassertion. This is a structural firewall built in to protect the margins of incumbent firms, ensure that innovations provide value, and build trust between actors. This firewall is bolstered by the parallel deployment strategy that recommends innovations are first given a trial run alongside existing incumbent workflows to reduce disruption, provide learnings to innovators, and build an evidence base for innovations. These structural features are built into the co-creation agreement and are therefore not dependent on goodwill. Thirdly, the compensation structures outlined in the framework are a mechanism that converts innovation commitment into a financially motivated decision, which operates on the exact logic that survival logic thrives – the incumbent will have a structurally induced financial incentive to protect the living lab under project stress.

The framework cannot, however, provide structural defence against catastrophic project stress. Incumbent survival instinct when facing insolvency will be to abandon all non-essential commitments to the project delivery, regardless of contractual or financial incentive. This is a genuine boundary condition of the framework. The living lab model is viable for firms operating under normal project stress but may not survive existential financial pressure. This is consistent with the micro-firm boundary condition acknowledged in Section 5.6.5. The framework's living lab model is designed for firms with sufficient organisational resilience to absorb innovation engagement under normal operating conditions, not for firms operating permanently at survival logic margins.

This boundary condition reflects a deliberate design choice rather than an oversight. Just as the framework recommends that innovators make a strategic decision about which segment of the bifurcated market to target first, the framework itself has made the same decision – prioritising mid-size firms with sufficient organisational resilience to engage with innovation under normal operating conditions. Attempting to design a framework applicable to every firm type simultaneously would reproduce the platform logic failure the framework explicitly critiques – trying to boil the ocean rather than solving one problem perfectly for a specific market segment. The micro-firm segment remains an important future research and policy challenge, but one that requires a different intervention logic than the CCF provides.

6.3.2 Stress Test Two: Parallel Deployment vs Observability

Although Chapter 2 identified significant limitations in Rogers' DOI framework – specifically its inability to account for technologies that do not stabilise before diffusion begins – the observability concept from his five adoption factors was not subject to that critique and remains analytically applicable here. The stress test applies this specific component rather than endorsing the broader diffusion model. Rogers' (2003) Diffusion of Innovations framework identified observability as a prerequisite for adoption. This explains that for an innovation to disseminate, it needs to be observable and the results must be visible. The framework's near-zero disruption parallel deployment strategy deliberately minimises the innovation's visibility to avoid competing with project delivery. The second stress test asks: if the AI operates invisibly alongside existing workflows, how will it generate observable results that the Asymmetric Adoption Threshold requires for full-deployment to be triggered?

The framework has three mechanisms that aim to ameliorate the issue of observability. Firstly, parallel deployment is inherently different from invisible deployment; in fact, the entire goal of parallel deployment is observability – being able to compare in real-time the conventional workflow with the AI equivalent. In this way, the parallel deployment strategy provides more observability because conventional adoption discards the real-time comparison. Real-time is key here because no two projects are exactly the same, so only comparing past project data is insufficient in evaluating an innovation. Secondly, the mechanism is the co-creation agreement's performance-metric emphasis. This structurally creates observability, where documented evidence generated to trigger full-deployment is the kind of information required in Rogers' concept of innovation diffusion. Thirdly, the final mechanism is that the human-in-the-loop protocol creates observability at the error level – every reviewed AI output produces a visible comparison between the AI recommendation and human judgment.

Observability of outcomes is higher in parallel deployment, but the observability of processes is lower due to the nature of AI technology. An incumbent can observe that the AI produced a better estimate or design, but without understanding why. Paradoxically, AI's reasoning can be opaque, and this opacity is a genuine Sim-to-Real gap that future research should address through transparency mechanisms that make AI reasoning legible to non-technical construction practitioners. This opacity paradox is itself a manifestation of the Moving Target Trap, where AI's improving capability outpaces the development of transparency mechanisms that would make that capability legible to non-technical users, meaning the observability gap may widen before it narrows. The framework addresses outcome observability but does not fully address process transparency.

6.3.3 Stress Test Three: Domain Expert Integrator vs Plasticity

Star and Griesemer's (1989) boundary object theory requires that boundary objects maintain sufficient plasticity to adapt to variable local conditions across different communities of practice. AI tools are characteristically brittle – dependent on specific data sets and, in the case of residential construction, are geographically locked by Regulatory Non-Replicability. This issue is exacerbated by Canada's extreme conditions of site-variability. The third stress test questions whether a domain expert integrator can serve as a proxy for the plasticity that AI technology structurally lacks.

The framework has two responses: RAG architecture and the domain expert integrator's role. The RAG architecture recommendation is explicitly a plasticity mechanism – by restricting the AI to a reconfigurable closed-loop domain-specific data pool, the framework acknowledges AI's brittleness and proposes a structural response. The RAG data pool can be reconfigured for different site conditions, jurisdictions, and project types without rebuilding the tool from scratch. The domain expert integrator's role is specifically designed to perform the translation function that compensates for AI rigidity. The expert bridges the gap between the AI's structured data requirements and the construction site's variable operational reality. The integrator does not make the AI more adaptable, but they perform the boundary-spanning function that manages the gap between AI rigidity and site variability.

The framework can falter if site variability exceeds the domain expert integrator's capacity to translate. This can happen if conditions are sufficiently novel that neither the RAG data pool nor the integrator's domain knowledge can account for them. In this case, the AI tool will fail. This maps onto the Regulatory Non-Replicability finding. Solutions validated in one site context do not automatically transfer to significantly different site contexts, even with a skilled integrator present. This is a genuine and irreducible boundary condition that the framework acknowledges but cannot eliminate. Extending the integrator's effective plasticity boundary requires the regulatory harmonisation pathway recommended in Section 5.2.3. Without cross-jurisdictional recognition frameworks, the integrator's domain knowledge remains as geographically bounded as the AI tool itself.

6.3.4 Stress Test Four: AI Errors, Liability, and Trust Erosion.

Orlikowski and Gash (1994) established that when different organisational groups hold incompatible technological frames, the result is systematic rather than accidental implementation failure. The framework's co-creation model requires innovators, incumbents, and integrators to develop sufficient frame alignment to function as collaborative partners. But when an AI tool produces an error with real physical or financial consequences on a construction site, the technological frames each actor carries will produce incompatible liability attributions – the innovator frames it as a data scope problem, the incumbent frames it as a product failure, and the legal-professional framework frames it as practitioner negligence. The final stress test asks: Does the framework provide sufficient frame alignment mechanisms to survive the interpretive crisis that an AI error on site would produce?

Three mechanisms in the framework partially address frame alignment. The domain expert integrator performs continuous frame translation. This role is needed to bridge the cognitive gap between construction's project logic and the innovator's product logic before a crisis occurs. The co-creation agreement's explicit contribution documentation creates a shared interpretive record, and both parties have agreed in writing what the tool is supposed to do, what data it draws from, and what human review protocols apply. This shared record provides a common reference point when frames diverge under pressure. The human-in-the-loop protocol distributes interpretive responsibility and reviewed outputs are professional judgments, not direct AI recommendations, which at minimum engages the legal-professional liability frame on terms it recognises.

These mechanisms reduce the technological frame gap under normal operating conditions but cannot eliminate it under crisis conditions. When physical harm or significant financial loss occurs, the institutional pressure to assign discrete blame will overrule the co-creation logic's responsibility model. The sociomaterial entanglement of human and AI decision-making means the output cannot be cleanly separated from the inputs both parties contributed. This resists the discrete causal chain that liability law requires, and this gap cannot be addressed through framework design alone. It requires the regulatory modernisation that Section 5.2.5's regulatory agility recommendation gestures toward but does not fully develop, specifically the establishment of professional standards for AI-assisted practice that distribute liability across the co-creation relationship rather than assigning it entirely to the licensed practitioner.

The most significant finding of this stress test is not that AI errors create legal problems; it is that they activate the full weight of technological frames simultaneously. This has the potential to destroy the collaborative relationship that the framework's three stages were designed to build. The interpretive crisis produced by an AI error may be more damaging to the interface than the error itself. This observation also reframes the Credibility Threshold finding from Chapter 4 – domain expertise is not just a trust prerequisite but a frame alignment prerequisite. An integrator with sufficient construction domain expertise provides the shared foundation that makes collaborative responses possible. Without that shared foundation, an AI error does not just fail a project; it confirms every incumbent's worst assumptions about innovation and raises the Credibility Threshold for every innovator who follows.

6.4 Net Transferability Assessment

The four stress tests collectively demonstrate that the CCF survives theoretical scrutiny with partial rather than complete answers to each challenge. The living lab survives normal project stress through phase-specific and equity agreements, but not existential financial pressure. Parallel deployment generates structured observability through real-time comparison and performance metrics, but not process transparency. The domain expert integrator maintains boundary-spanning within their domain knowledge but not beyond it. And the co-creation model reduces technological frame divergence under normal operating conditions, but cannot eliminate the interpretive crisis that an AI error would produce under pressure.

In each case, the framework has a structural response that reduces but does not eliminate the Sim-to-Real gap. This is not a failure of framework design, but it reflects the honest limits of what any conceptual framework can prescribe in the face of structural conditions as deeply entrenched as those documented in Chapters 2 and 4. A framework that claims to eliminate these gaps entirely would be analytically dishonest.

The framework's most significant transferability challenge emerges not from any single stress test but from their compounding interaction. A high-stress project environment with variable site conditions, an underfunded or absent integrator, and a recent AI error activates all four stress test vulnerabilities simultaneously. Thus, survival logic abandons the living lab, site variability exceeds the RAG architecture's plasticity, parallel deployment's observability advantage is undermined by the integrator's reduced capacity to facilitate comparison, and incompatible technological frames produce an interpretive crisis that the co-creation relationship cannot survive. This compounding scenario represents the framework's outer boundary condition, and where the specific combination of circumstances under which its transferability breaks down most completely. This boundary condition defines not just where the framework fails but why it fails and what conditions would need to change for it to succeed. The enabling layer's policy recommendations are not supplementary additions to the framework, but exist as the preconditions that prevent the compounding failure scenario from becoming the default outcome. Without policy intervention to socialise first-mover risk, fund embedded integrators, harmonise regulatory frameworks, and modernise professional liability standards, the compounding failure scenario is not an edge case. It is the structural baseline from which the framework is trying to depart. The net assessment is that the CCF is analytically credible, empirically grounded, and practically motivated. This framework is a theory-building contribution that meets the rigorous standard of conceptual framework development while generating specific and testable propositions for future empirical confirmation.

6.5 Implications for Future Research

The four stress tests applied in Section 6.3 each produced residual gaps that the framework acknowledges but cannot resolve through design alone. These gaps generate specific empirical questions that future research should address to extend the framework's reach and validate its transferability beyond the theory-building scope of this thesis.

Q1: Emerging from Stress Test One's survival logic vulnerability: to what extent can co-creation agreements, financial stakes, and AI error preventative measures systematically stabilise tentative relationships at the interface when faced with project stress?

Q2: Emerging from Stress Test Three's plasticity boundary: Under what conditions of site variability does the domain expert integrator's boundary-spanning capacity fail, and what augmentation mechanisms can be identified to extend that capacity? E.g. additional domain expertise, expanded RAG architecture, or AI-specific training.

Q3: What minimum enabling layer investment in government programme scale, capital instrument design, and regulatory harmonisation scope is required to prevent the compounding failure scenario identified in Section 6.4 from becoming the structural baseline at the interface rather than the exceptional edge case?

Q4: What professional liability frameworks for AI-assisted construction practice would adequately distribute responsibility across co-creation relationships? Which institutional actors, licensing bodies, industry associations, or government regulators are best positioned to develop and enforce them?

7. Conclusion:

This chapter presents the conclusions of a study into the barriers and enablers determining whether entrepreneurial AI innovations successfully navigate the commercialisation interface with Canadian residential construction firms. The findings show that commercialisation failure in this context is not primarily a technology or an attitudinal problem, but a structural and relational problem. The resolution of this problem requires deliberate reorientation of the commercialisation encounter from competition logic to interaction logic. This must be enabled by institutional infrastructure that the market cannot generate unilaterally. The chapter proceeds through a direct answer to the research question, a statement of theoretical and practical contributions, an acknowledgement of limitations, directions for future research, and a closing statement.

7.1 Answering the Research Question

7.1.1 Barriers and Enablers to AI Adoption at the Interface

The primary research question of this thesis focused on the barriers and enablers to AI adoption at the interface. The answer to this question is architectural rather than additive. The barriers that determine a successful entrepreneurial interface navigation are not independent conditions that could be addressed sequentially. Instead, they are mutually reinforcing dimensions of a structural void that collectively diminish the probability that the conditions of a productive encounter are present. Five primary barrier clusters were identified and documented. The frozen industry, described as thirty years of normalised stagnation reinforced by the Real Estate Appreciation Shield, effectively removes the urgency that innovation adoption requires. The hostile interface makes the commercialisation encounter systematically difficult regardless of product quality. This interface is made hostile by the Credibility Threshold, the no-first-mover dynamic, the hallucination trust deficit, and the structural conflict of interest. Structural incapacity, characterised by Operational Capture, data infrastructure failure, Regulatory Non-Replicability, and ecosystem coordination failure, prevents even willing actors from engaging productively. Finally, the enabling conditions identified empirically – problem-first orientation, domain expertise, open-minded incumbent co-creation, and the emerging systems integrator role – are present but rare, fragile, and dependent on individual relationships rather than systemic structuresThese conditions are present but rare, fragile, and dependent on individual relationships rather than systemic structures. The primary determinant of successful interface navigation is not the innovation's technical quality but the conditions of encounter – specifically, whether the Credibility Threshold is cleared, whether co-creation conditions exist, and whether the enabling institutional infrastructure reduces the structural barriers that individual actors cannot overcome unilaterally.

7.1.2 Necessary Conditions for Adoption

For the second part of the primary research question, the necessary conditions for adoption to begin were explored. Three conditions were found necessary to be simultaneously present. First, the Credibility Threshold must be cleared, where the innovator possesses or has access to sufficient construction domain expertise to perform the integrating function that the industry cannot generate internally. Second, the incumbent must arrive at the interface as a co-creator rather than a passive evaluator. The incumbent must arrive with articulated pain points, protected innovation bandwidth, and organisational commitment to the living lab principle. Third, the enabling institutional infrastructure must exist to make these conditions reproducible rather than dependent on rare individual relationships. When all three conditions are met, adoption simultaneously becomes viable. When anyone is absent, the interface reverts to its default hostile state. These three conditions form the empirical foundation of the CCF developed in Chapter 5, and each layer of the framework is designed to make one or more of these conditions systematically achievable rather than accidentally rare.

7.2 Theoretical Contributions

7.2.1 The Commercialisation Interface as Unit of Analysis

Existing commercialisation frameworks locate their unit of analysis either at the technology level or at the firm level. These two approaches examine innovation characteristics that predict adoption and organisational capabilities that enable adoption. This thesis introduces the meso-level commercialisation interface as a distinct unit of analysis with its own structural logic, institutional dynamics, and failure modes. This reframe reveals barriers and enablers that technology-level and firm-level analyses cannot see. The encounter conditions that determine whether a productive engagement is possible, regardless of product quality or organisational capability is only observable at the meso-level. The interaction interface is positioned as a dynamic threshold-dependent variant of Star and Griesemer's boundary object concept – one that can fail entirely if the Credibility Threshold is not cleared, a condition the original boundary object theory does not anticipate. By establishing the interaction interface as the primary unit of analysis, this thesis opens a new analytical space for commercialisation research. One that treats the encounter itself as the object of intervention rather than the technology being commercialised or the firm being asked to adopt it.

7.2.2 The Credibility Threshold as Prerequisite Condition

The thesis establishes that domain expertise in the innovator is a prerequisite condition for interface engagement rather than a competitive advantage. Below this threshold, the interface is non-functional regardless of product quality, ROI projections, or technical superiority. This finding addresses the blind spot in Slaughter's innovation typology, which accounts comprehensively for product characteristics but contains no variable for innovator characteristics. Effectively, this thesis extends Slaughter's framework into a new explanatory mechanism. The Credibility Threshold was empirically grounded in the Grey Hair Effect: a practitioner-generated concept that captures the cognitive offloading mechanism through which bandwidth-constrained incumbents transfer the evaluation function to domain-expert innovators. This mechanism reframes domain expertise from a legitimacy signal to a structural functional necessity. This reframe has direct implications for how commercialisation support programmes, accelerators, and policymakers should evaluate and select ConTech startups. Domain expertise is not a nice-to-have characteristic, but the prerequisite condition that determines whether public or private investment in commercialisation support will produce viable interface encounters.

7.2.3 The Moving Target Trap

The Moving Target Trap names a paradox specific to rapidly evolving technologies, where rational caution produces irrational strategic outcomes because the stability being waited for structurally cannot arrive. Rogers' DOI framework predicts that innovations diffuse as they stabilise and demonstrate relative advantage. AI's developmental trajectory violates this assumption, where each iteration changes the technology's capability profile before incumbents can evaluate the previous one. Firms waiting for AI to stabilise are waiting for a condition that AI's development logic structurally prevents from arriving, while their capability gap widens and their organisational fluency fails to develop. This paradox extends DOI theory beyond its acknowledged limitations into a new failure mode applicable to any technology whose development pace exceeds institutional adoption capacity.

7.2.4 Regulatory Non-Replicability

Regulatory Non-Replicability synthesises Khanna and Palepu's institutional voids framework, Szulanski's replication theory, and Arthur's path dependency argument into a unified explanation for why AI solutions remain locally rather than nationally effective. Solutions validated in one Canadian province cannot be deployed elsewhere without significant redesign or recertification. Structurally preventing scalability in fragmented multi-jurisdictional markets, which effectively resets the commercialisation clock and prevents the compound scaling that drives AI improvement in less fragmented sectors. This concept was empirically confirmed by multiple participants including a policymaker with direct housing policy experience and has implications beyond Canadian construction for any sector operating across heterogeneous regulatory environments.

These theoretical contributions are complemented by two contextually specific empirical findings. First, Operational Capture, which reframes incumbent non-adoption from cultural resistance to structural bandwidth poverty. Second, the Real Estate Appreciation Shield, which identifies a market distortion that suspended the productivity pressures in Canada that would otherwise have forced innovation adoption. Both contributions extend the theoretical analysis into mechanisms that the existing literature had not previously named.

7.3 Practical Contributions

7.3.1 For Innovators

The thesis recommends the Embedded Value Escalation Strategy. A three-stage commercialisation approach that sequences market engagement to clear the Credibility Threshold before attempting full deployment. Stage One requires industry logic adoption, Stage Two requires near-zero disruption parallel deployment, and Stage Three requires controlled full deployment with deliberate institutionalisation mechanisms that prevent successful pilots from dissolving when project teams disperse. The niche-first entry strategy this produces is not a permanent commercial limitation; instead, it is the sequenced pathway through which trust, domain fit, and validated evidence are built before broader digital transformation engagement becomes viable.

7.3.2 For Incumbents

The thesis recommends a three-step reposition from sceptic to co-creator. Domain expertise should be reframed from an evaluative standard to an innovation resource. Incumbent operational knowledge is the input that makes AI tools viable in construction conditions, and pain points should be articulated before approaching the interface. Incumbents should be prepared with the ability to answer the question "what would I pay to have solved right now?" both as a filter for evaluating innovations and the brief that guides productive co-creation. The living lab principle should be adopted as an organisational commitment where offering site access and operational knowledge in exchange for equity or co-development credit aligns incentives structurally and transforms the interface from a transactional encounter into a co-development relationship.

7.3.3 For Policymakers

The thesis recommends four enabling institutional infrastructure investments. First, a government-funded embedded innovator partnership programme that places domain-expert innovators within incumbent firms as systems integrators. This simultaneously addresses the missing integrator vacancy, incumbent bandwidth poverty, and the no-first-mover coordination failure. Second, construction-specific innovation infrastructure that provides the physical facilities ConTech startups require. Third, a national regulatory harmonisation pathway that allows solutions validated in one province to be recognised in others. And fourth, ConTech-specific capital instruments with return horizons aligned with construction's adoption pace rather than VC growth logic.

7.4 Limitations

This thesis has four main limitations: The simultaneous behaviour change requirement, the scaling paradox, the sample and context boundary, and the Moving Target Trap's recursive application. The framework requires coordinated change from multiple actors that market mechanisms alone cannot produce, and these changes conflict with their dominant institutional logics. The niche-first strategy conflicts with the VC growth logic that motivates innovation in ways the framework partially but not fully resolves. Twelve interviews from

Canadian residential construction constitute a theory-building contribution whose propositions require further empirical confirmation rather than proven prescriptions applicable across all construction contexts. The framework's specific recommendations are temporally bound by AI technologies' 2025 capabilities and require periodic reassessment as the technology evolves. Taken together, these limitations do not diminish the thesis's analytical contribution, but they define its scope. Further, they generate the specific empirical questions that future research must address to extend the CCF from a theory-building proposition into an empirically validated intervention model.

7.5 Future Research Directions

7.5.1 AI Process Transparency in Construction

The opacity paradox identified in Stress Test Two, where incumbents can observe AI outcomes but not AI reasoning, suggests a research agenda around transparency mechanisms that make AI decision logic legible to non-technical construction practitioners. Future research should examine which transparency formats are most effective at building incumbent risk calibration in high-stakes construction environments.

7.5.2 Integrator Capacity and RAG Architecture Limits

The plasticity boundary identified in Stress Test Three suggests a research agenda around the conditions under which domain expert integrators fail and the augmentation mechanisms that extend their capacity. Future research should examine the relationship between site variability, RAG data pool specificity, and integrator domain knowledge depth. This research could be mapping the boundary conditions under which the framework's plasticity mechanisms remain effective.

7.5.3 Professional Liability Frameworks for AI-Assisted Construction

The interpretive crisis identified in Stress Test Four suggests a research agenda around professional liability for AI-assisted construction practice. Future research should examine how licensing bodies, industry associations, and government regulators can develop liability frameworks that distribute responsibility across co-creation relationships without undermining the collaborative conditions that make co-creation viable.

7.5.4 Enabling Layer Minimum Viable Investment

The compounding failure scenario identified in Section 6.4 suggests a research agenda around the minimum enabling layer investment required to prevent the compounding failure from becoming the structural baseline. Future research should examine which enabling layer components produce the greatest reduction in compounding failure risk per unit of policy investment. This could explore partnership programmes, capital instruments, regulatory harmonisation, or physical infrastructure.

7.6 Closing Statement

Nothing about building houses has changed in thirty years; some practitioners still track multimillion-dollar projects in written notebooks. This observation, offered by an industry incumbent, captures the empirical reality documented in this research: a sector where stagnation is not coincidental, but architecturally produced. The Canadian residential construction sector is governed by a system of structural conditions and institutional logic clashes that were never designed to accommodate externally driven AI innovation. As this thesis has demonstrated, no amount of superior technology or persuasive pitching can overcome this systemic inertia without deliberate structural intervention.

By reframing the industry's productivity lag from a technology adoption failure into a structural and relational crisis, this research identifies the precise site where intervention must occur. The Interaction Interface is the Valley of Death where structural voids, conflicting logics, and innovator characteristics collide. While this interface is currently Architecturally Hostile, the conditions required to make it productive are knowable, specifiable, and achievable. However, these conditions require deliberate, collaborative design rather than a reliance on passive market emergence.

The domain-expert innovator navigating this interface is performing a function the industry cannot currently generate internally: bridging incompatible logics, offloading cognitive burdens, and producing the domain fit necessary for AI to survive on-site. When this role is embedded through the CCF, the interface transitions from a site of conflict to a site of transformation. Ultimately, this framework delivers the roadmap for the relational and structural mending that Canadian residential construction's AI future requires.

References

Adner, R. (2006): Match your innovation strategy to your innovation ecosystem. In: Harvard Business Review, vol. 84 (2006), pp. 98-107.

Arthur, W. B. (1989): Competing technologies, increasing returns, and lock-in by historical events. In: The Economic Journal, vol. 99 (1989), pp. 116-131.

Bossink, B. A. G. (2004): Managing drivers of innovation in construction networks. In: Journal of Construction Engineering and Management, vol. 130 (2004), pp. 337-345.

Bridgit (2023): Tech fatigue in construction. URL: https://gobridgit.com/blog/tech-fatigue-construction/ (accessed 24.03.2026).

Business Development Bank of Canada (BDC) (2022): Meet the Canadian innovators behind construction’s digital coming of age. n.p.: BDC.

Crofton, J.-P. (2024): What's wrong with the diffusion of innovation theory?. n.p.: Edge of Possible Consultancy.

Dubois, A./Gadde, L.-E. (2002): The construction industry as a loosely coupled system: Implications for productivity and innovation. In: Construction Management and Economics, vol. 20 (2002), pp. 621-631.

Fichman, R. G. (2004): Going beyond the dominant paradigm for information technology innovation research: Emerging concepts and methods. In: Journal of the Association for Information Systems, vol. 5 (2004), pp. 314-355.

Friedland, R./Alford, R. R. (1991): Bringing society back in: Symbols, practices, and institutional contradictions. In: Powell, W. W./DiMaggio, P. J. (ed.): The new institutionalism in organizational analysis. n.p.: University of Chicago Press.

Gann, D. M./Salter, A. J. (2000): Innovation in project-based, service-enhanced firms: The construction of complex products and systems. In: Research Policy, vol. 29 (2000), pp. 955-972.

Ghimire, P./Kim, K./Acharya, M. (2024): Opportunities and challenges of Generative AI in construction industry: Focusing on adoption of text-based models. In: Buildings, vol. 14 (2024), pp. 220.

Havenvid, M. I. (2015): Competition versus interaction as a way to promote innovation in the construction industry. In: The IMP Journal, vol. 9 (2015), pp. 46-63.

Hobday, M. (2000): The project-based organisation: An ideal form for managing complex products and systems? In: Research Policy, vol. 29 (2000), pp. 871-893.

Innovation, Science and Economic Development Canada. (2024). *Residential building construction - 2361: Employer establishments by employment size category and province/territory*. Government of Canada.

Khanna, T./Palepu, K. G. (2010): Winning in emerging markets: A road map for strategy and execution. n.p.: Harvard Business Press.

Koebel, C. T. et al. (2003): The diffusion of innovation in the residential building industry. n.p.: U.S. Department of Housing and Urban Development.

Koeleman, J. et al. (2019): Decoding digital transformation in construction. URL: https://www.mckinsey.com/business-functions/operations/our-insights/decoding-digital-transformation-in-construction (accessed 24.03.2026).

Kurup, P./Pless, S./Grosskopf, K. (2024): Understanding the challenges of financing modular construction: A case study for prospective multifamily units. n.p.: National Renewable Energy Laboratory.

Kussl, S./Wald, A. (2025): The role of construction clients in digital innovation: Insights from scenario analysis. In: Construction Management and Economics, vol. 43 (2025), pp. 381-404.

Laberge, M. (2025): Unlocking housing supply: Why scale really matters. URL: https://www.cmhc-schl.gc.ca/observer/2025/unlocking-housing-supply-why-scale-really-matters (accessed 24.03.2026).

Latour, B. (2005): Reassembling the social: An introduction to actor-network-theory. n.p.: Oxford University Press.

Law, J. (1992): Notes on the theory of the actor-network: Ordering, strategy, and heterogeneity. In: Systems Practice, vol. 5 (1992), pp. 379-393.

Leonardi, P. M./Barley, S. R. (2010): What’s under construction here? Social action, materiality, and power in constructivist studies of technology and organizing. In: Academy of Management Annals, vol. 4 (2010), pp. 1-51.

ab Iorwerth, A. (2025). *Framework for change: Productivity in housing construction*. Canada Mortgage and Housing Corporation (CMHC)

Lyytinen, K./Damsgaard, J. (2001): What's wrong with the diffusion of innovation theory? The case of a complex and networked technology. In: anon. (ed.): Working Conference on Diffusing Software Product and Process Innovations. n.p.: Springer.

Markham, S. K. (2002): Moving technologies from lab to market. In: Research-Technology Management, vol. 45 (2002), pp. 31-42.

McCoy, A. P. et al. (2008): Understanding the effect of concurrent commercialization practices on product adoption thresholds. In: COBRA: The RICS Construction and Building Research Conference 2008. Dublin: n.p.

McCoy, A. P./Thabet, W. T./Badinelli, R. D. (2008): Towards establishing a domain-specific commercialization model for innovation in residential construction. In: Construction Innovation, vol. 8 (2008), pp. 137-155.

Meng, X./Brown, A. (2018): Innovation in construction firms of different sizes: Drivers and strategies. In: Engineering, Construction and Architectural Management, vol. 25 (2018), pp. 1210-1225.

Mohammed, S. (2023): Going beyond the curve: Understanding the limitations of the diffusion of innovation model. n.p.: Medium.

Onetrace (2026): How AI is transforming construction scheduling in 2026. URL: https://onetrace.com/journal/ai-construction-scheduling (accessed 24.03.2026).

Orlikowski, W. J. (2007): Sociomaterial practices: Exploring technology at work. In: Organization Studies, vol. 28 (2007), pp. 1435-1448.

Orlikowski, W. J./Gash, D. C. (1994): Technological frames: Making sense of information technology in organizations. In: ACM Transactions on Information Systems (TOIS), vol. 12 (1994), pp. 174-207.

Ozorhon, B./Abbott, C./Aouad, G. (2010): Construction innovation: Towards a project-level analysis. In: Proceedings of the 26th Annual ARCOM Conference. n.p.: n.p.

Papadonikolaki, E./Krystallis, I./Morgan, B. (2020): Digital transformation in construction: Systematic literature review of evolving concepts. In: Engineering Project Organization Conference (EPOC). n.p.: n.p.

Pfeffer, J./Salancik, G. R. (1978): The external control of organizations: A resource dependence perspective. n.p.: Harper & Row.

Prencipe, A./Tell, F. (2001): Inter-project learning: Processes and outcomes of knowledge codification in project-based firms. In: Research Policy, vol. 30 (2001), pp. 1373-1394.

RBC (2024): Building blocks: Addressing Canada’s housing crisis with innovation and prefabricated construction. n.p.: Royal Bank of Canada.

Rogers, E. M. (2003): Diffusion of innovations. 5th ed. n.p.: Free Press.

Samuelson, O./Stehn, L. (2023): Digital transformation in construction – a review. In: Journal of Information Technology in Construction (ITcon), vol. 28 (2023), pp. 385-404.

Saqib, S. I./Allen, M. M. C. (2024): Institutional logics in play at work: How applying an institutional logics approach to employees’ intentions to quit contextualizes HRM. In: The International Journal of Human Resource Management, vol. 35 (2024), pp. 2839-2862.

Schad, J. et al. (2016): Paradox research in management science: Looking back to move forward. In: The Academy of Management Annals, vol. 10 (2016), pp. 5-64.

Schildt, H. (2022): The institutional logic of digitalization. In: anon. (ed.): Research in the Sociology of Organizations. n.p.: Emerald Publishing Limited.

Searle, B./Rankin, J. H. (2025): How innovation is impacting productivity in the construction industry: A review and case study of the Canadian research landscape. In: Proceedings of the 23rd CIB World Building Congress. West Lafayette, IN: n.p.

Slaughter, S. E. (1998): Models of construction innovation. In: Journal of Construction Engineering and Management, vol. 124 (1998), pp. 226-231.

Smith, W. K./Lewis, M. W. (2011): Toward a theory of paradox: A dynamic equilibrium model of organizing. In: Academy of Management Review, vol. 36 (2011), pp. 381-403.

Spare, N./Menard, A./Snyder, J. P. (2018): Big data: Big questions for the E&C industry. n.p.: FMI.

Star, S. L./Griesemer, J. R. (1989): Institutional ecology, 'translations' and boundary objects: Amateurs and professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39. In: Social Studies of Science, vol. 19 (1989), pp. 387-420.

Statistics Canada (2024): The construction sector is not operating at full capacity—here are some data that could explain why. URL: https://www.statcan.gc.ca/o1/en/plus/6596-construction-sector-not-operating-full-capacity-here-are-some-data-could-explain-why (accessed 24.03.2026).

Thornton, P. H./Ocasio, W. (2008): Institutional logics. In: Greenwood, R. et al. (ed.): The SAGE handbook of organizational institutionalism. n.p.: SAGE Publications.

Walzer, A. N. (2025): Assessing multi-level uncertainties in construction robotics entrepreneurship by means of methodological pluralism. URL: https://doi.org/10.3929/ethz-b-000705101 (accessed 24.03.2026).

Weick, K. E. (1976): Educational organizations as loosely coupled systems. In: Administrative Science Quarterly, vol. 21 (1976), pp. 1-19.

Williamson, O. E. (1994): Transaction cost economics and organization theory. In: Smelser, N. J./Swedberg, R. (ed.): The handbook of economic sociology. n.p.: Princeton University Press.

Winch, G. M. (2003): How innovative is construction? Comparing aggregated data on construction innovation and other sectors—a case of apples and pears. In: Construction Management and Economics, vol. 21 (2003), pp. 651-654.

Winter, S. G./Szulanski, G. (2001): Replication as strategy. In: Organization Science, vol. 12 (2001), pp. 730-743.

Yang, L. et al. (2025): Achieving on-site trustworthy AI implementation in the construction industry: A framework across the AI lifecycle. In: Buildings, vol. 15 (2025), pp. 21.

Yin, R. K. (2018): Case study research and applications: Design and methods. 6th ed. n.p.: SAGE Publications.

Zachariah, J./Goldsmith, T. (2022): Laying foundations: Technological maturity in Canada’s construction sector. n.p.: Brookfield Institute for Innovation + Entrepreneurship.