

Form Follows Forecast:
Monte Carlo Based Feasibility Analysis of Architectural Partis

Nicholas Miranda

A thesis
submitted in partial fulfillment of the
requirements for the degree of

Master of Architecture

University of Washington
2026

Committee:
Narjes Abbasabadi
Arthur Acolin

Program Authorized to Offer Degree:
Architecture

Abstract
University of Washington

Form Follows Forecast:
Monte Carlo Based Feasibility Analysis of Architectural Partis

Nicholas Miranda

Chair of the Supervisory Committee:
Narjes Abbasabadi
Department of Architecture


Contemporary American multifamily development is typically evaluated through financial models that reduce architecture to scalar inputs: gross floor area, story count, cost per square foot, rent, and exit value, among others. These variables are necessary, but they are largely blind to geometry. This thesis proposes an abstracted proxy framework for making the architectural parti more financially legible by translating early massing decisions into two composite indices: Architectural Complexity C and Asset Quality Q . These indices connect measurable design conditions to cost- and value-facing consequences, then enter a Monte Carlo pro forma that samples market and delivery uncertainty. The framework is tested on a hypothetical market-rate multifamily tower in Seattle's Eastlake neighborhood through a sequence of authored massing iterations. The results show that spatial ambition and financial robustness do not automatically align: balconies, rounded forms, amenity moves, and added height can improve architectural value while weakening underwriting performance. Rather than a finished predictive model, this thesis is a focused exploration of how architectural form might be quantified earlier, while the parti is still flexible enough to be tested, revised, and improved.

Contents

1	Introduction	1
1.1	The Wager	2
1.2	Define <i>Parti</i>	3
1.3	The Question	3
1.4	What Follows	4
1.5	A Note on What This Thesis Is Not	5
2	Background & Context	6
2.1	Probabilistic Evaluation in Real Estate and Development	6
2.2	Computational Design and Deterministic Economics	9
2.3	Uncertainty in Construction Delivery	9
2.4	Set-Based Reasoning as Methodological Precedent	11
2.5	The Missing Integration	11
3	Methodology	13
3.1	Financial Model	14
3.2	3D Experiment Space	16
3.3	Architectural Complexity (C)	18
3.4	Asset Quality (Q)	19
3.5	The Proxies	19
3.6	Normalization and Composite Scoring	26
3.7	How Height Enters the Model	28
3.8	Monte Carlo Simulation	32
3.9	From Distributions to Decisions	33
3.10	The Analytical Outputs	34
3.11	Transition-Based Diagnostic Analysis	36
4	Project Description and Model Inputs	38
4.1	Input Workbook Structure	38
4.2	Site and Planning Context	39
4.3	Architectural Scenario	41
4.4	Construction Cost Inputs	42
4.5	Revenue and Tenure Inputs	43
4.6	Financing Inputs and Robustness Thresholds	45
4.7	Stochastic Inputs	46
4.8	Design Variables Reserved for the Experiment	47
5	The Experiment	48
5.1	Iteration 1 – Baseline	50
5.2	Iteration 2 – Perimeter Setback Balconies	54
5.3	Iteration 3 – Swept Southwest Corner	58
5.4	Iteration 4 – Rooftop Amenity Deck	63
5.5	Iteration 5 – Height Extrusion	68

5.6	Trajectory Summary	73
5.7	Tenure Test: Apartment Baseline Against Condominiums	76
6	Conclusion: Flying the Envelope	80
6.1	What the Demonstrator Actually Showed	80
6.2	What This Thesis Contributes	82
6.3	Limitations of the First Pass	82
6.4	Disciplinary Agency After the Pro Forma	83
6.5	The Demonstrator	84
	Bibliography	85
7	Appendix	95
7.1	Variable and Symbol Index	95

Note: Generative AI tools, including ChatGPT, Claude, Gemini, and xAI, were used during the preparation of this thesis for brainstorming, drafting assistance, copyediting, code debugging, and formatting support. All analytical claims, citations, model assumptions, calculations, and final editorial decisions remain the responsibility of me. As such, AI-generated text was revised, verified, and integrated by me rather than used as an independent scholarly source.



1 Introduction

On 28 January 2025, an experimental aircraft called XB-1 climbed to thirty-five thousand feet above the Mojave Desert and broke the sound barrier [Boom Supersonic, 2025]. It did so quietly, in two senses: it drew little notice beyond the aerospace press, and the sonic boom never reached the ground. The aircraft was a one-third-scale demonstrator built by Boom Supersonic, a startup founded on the premise that commercial supersonic flight, abandoned in 2003 when British Airways and Air France retired the Concorde, was not a closed problem. Conventional wisdom held that the Concorde had proven the category unviable. Boom's wager was that the wisdom had mistaken a *historical* failure for a *physical* one. What killed the Concorde, on this reading, was not the aerodynamics of flight above Mach 1 but the methodology of 1960s aerospace applied to 1960s cost structures: bespoke titanium fabrication, analog flight controls, and a financial model that closed only with state subsidy. Sixty years of advances in composite materials and computational design had never been allowed to revisit the question. The category had been declared dead before the tools that might revive it existed.

A version of this story is now being told inside American multifamily development, though it is usually framed incorrectly. The high-budget building, conceived in the twentieth century as a civic and cultural beacon, has not vanished from the American landscape: supertall residential towers still rise, flagship institutions still commission signature architects, and trophy mixed-use developments still anchor downtown skylines. What has changed is not the presence of architectural ambition but the conditions under which it can be realized. The relationship between design intent and financial logic has tightened to the point that most projects are resolved less as expressions of architectural possibility than as outcomes of economic constraint.

The buildings remain expensive, but the expense no longer buys latitude. Formal moves such as cantilevers, articulated envelopes, and material differentiation are not merely aesthetic; they are economically consequential, judged by their effect on cost, risk, and return. Architecture has not been displaced by finance so much as absorbed into it. Design decisions pass through feasibility criteria that sit upstream of most architectural differentiation, and the range of pursuable outcomes narrows accordingly. The result is a built environment that can look expressive while staying structurally conservative. Projects that do depart from the standard configuration tend to do so under a specific financial logic, as high-end globally marketed assets or as heavily constrained efficiency-driven developments, rather than as open explorations of form. The architecture has not disappeared; the space in which it can vary has.

This is not universal, and it is not best read as a failure of design culture. Cities such as Toronto, Melbourne, Amsterdam, Istanbul, and Singapore continue to deliver mid- and high-rise residential that is formally considered, materially specific, and financially viable at scales the United States treats as structurally hard to reach. The point is not that these cities have solved design. It is that their development economics let a wider range of outcomes stay viable in the first place: lower rel-

ative construction costs, steadier financing, and more predictable regulation widen the feasible set, even if the resulting form tends toward a consistent register. What looks like a divergence in architectural quality is better understood as a divergence in constraint.

The American condition, then, is not a disciplinary failure but a systemic one. High construction costs, volatile inputs, entitlement friction, code, and conservative underwriting compress the space of viable outcomes before design can differentiate within it. The result is not the absence of ambition but its filtration: many possible buildings exist at the outset of a project, and few survive the economic criteria that decide what gets built. The architecture that reaches completion expresses less what architects can produce than what the system permits.

If the methodology exists elsewhere, the American condition is not terminal, though this thesis does not address it at the level of root causes. The recovery argued here does not come from resolving construction cost, reforming zoning, or restructuring capital markets, all of which a full correction would require. It comes from operating more precisely within the constraints those conditions impose. The financialized American pro forma has not removed the architect; it has narrowed the terms on which architecture can be evaluated. The building enters that model as a scalar: square footage, story count, cost per foot. It does not see geometry. It does not see articulation. It does not see the difference between a podium that anchors a street wall and a podium that defaults to it. The language in which feasibility is established has no native terms for the qualities architects are trained to produce, so those qualities enter as costs to be justified rather than values to be priced.

1.1 The Wager

The wager of this thesis is that the conversation between architecture and finance can be restructured by giving them a shared language. The tools to make architectural ambition financially legible already exist in adjacent fields; they have not yet been brought into architectural decision-making at the moment they matter most. What gets framed as a conflict of values is, in practice, a failure of translation under constraint: a pro forma that cannot encode geometry will read geometry as cost, and a financial model that cannot register quality will treat it as a luxury. If architecture is to shape the city rather than merely dress the assets that occupy it, the task is to build that encoding. The instrument developed here is the Safe Design Envelope, built to evaluate the *parti*, the early organizing idea of a project, whether the tower cantilevers, the plates taper, the core shifts, or the massing steps. This is the moment when form takes shape and its consequences are set, and also the moment the conventional pro forma is least equipped to read it, since the building still registers as an undifferentiated volume that can only be priced after the fact. By then the implications have already hardened into cost, schedule, and risk, leaving little room for revision. The premise of *Form Follows Forecast* is a simple inversion: not that form should follow finance, but that form-giving can be evaluated probabilistically at the moment it is conceived, while it can still change.

The Safe Design Envelope makes that evaluation operable. It encodes the *parti*

through two empirical indices, Architectural Complexity (C) and Asset Quality (Q), extracted from a simple Rhino model together with a small set of manual inputs that register architectural intent geometry cannot carry. The framework samples across the bounded space of plausible massings on a given site and evaluates each against tens of thousands of stochastic realizations of the market environment, returning a map of where in the design space financial robustness lives. The map does not prescribe a building, nor optimize toward a single correct answer. It is a flight envelope, not a flight plan: a bounded region within which the parti can be pursued with quantified confidence about its financial behavior, leaving the architect to fly the aircraft.

The methodological precedents are developed in Chapter 3 but warrant brief naming. The probabilistic treatment of development feasibility runs from Hoesli et al. [2006] through Loizou and French [2012] and, most directly, Leung [2014], who showed that Monte Carlo treatment of pro forma inputs yields qualitatively different strategic conclusions than the deterministic spreadsheet it replaces. The demonstration that architectural form is itself a priced variable came from Rong et al. [2020], whose hedonic regression across 3,037 Manhattan commercial transactions found pricing differentials of roughly $\pm 10\text{--}12\%$ attributable to specific external form features, independent of awarded-architect effects. These two precedents bound the gap this thesis fills: Leung treated the building probabilistically but as a scalar; Rong treated form as financially consequential but only *ex post*, after the building was built and transacted. Neither integrated the two into an instrument operable at the moment of the parti.

1.2 Define *Parti*

The term recurs throughout this thesis and deserves a precise definition. *Parti* (from the French *prendre parti*, to take a position) is the central organizing idea of a design, the basic scheme from which the rest of the building develops: the decision about how mass, structure, circulation, and program are fundamentally arranged. It is not a style or a motif but a commitment, the first move that conditions every subsequent one. Whether the tower cantilevers or rises straight, whether the plates taper or stack uniformly, whether the core centers or shifts, these are parti-level decisions. They are made early, on limited information, and they are expensive to reverse, which is exactly why they are the decisions this thesis sets out to evaluate while they remain open.

1.3 The Question

The question this thesis sets out to answer is not whether ambitious architecture can be financially viable, which is obviously yes in some cases and no in others, nor whether it should be, which is a question of values the discipline has already settled. It is narrower and methodological, and it operates in two registers.

At the methodological level: can architectural complexity and asset quality, the two design-state variables most consequential for financial performance, be en-

coded with enough fidelity to enter a probabilistic financial model alongside conventional inputs like cost, rent, financing, and exit assumptions? If they can, a second question follows: can such a model produce a feasibility instrument that distinguishes robust from vulnerable design states across a wide range of market conditions, in a way that informs architectural decision-making rather than displacing it?

At the disciplinary level the question shifts from possibility to consequence. If the instrument exists, what does it change? Does a tool that makes both the cost of ambition and the value of quality visible restore the architect to a meaningful position in the financial conversation, one that has been eroding for decades? And does it answer a more stubborn worry: whether a formally ambitious parti, on a constrained urban parcel, can satisfy architectural intent and financial defensibility at once, or whether the two are fundamentally at odds, as recent American practice tends to assume?

The answer, developed in Chapter 3 and applied in Chapter 5, is that the two conditions are not structurally incompatible. They are in tension, and that tension is what the Safe Design Envelope is built to make visible. The thesis does not claim better architecture will always pencil. It offers something more practical: when better architecture does pencil, the conditions that make it possible can be identified in advance; when it does not, the gap between ambition and feasibility can be measured, named, and addressed through design revision, program adjustment, or an explicit recognition that additional support is required, rather than quietly absorbed into a value-engineering process the architect barely steers.

1.4 What Follows

The remainder of the thesis develops the argument across five further chapters. Chapter 2 situates the work in the lineages it draws on, the financialization of architectural production, the probabilistic turn in real estate modeling, and the empirical literature on form as a priced variable, and identifies the gap the Safe Design Envelope is meant to occupy. Chapter 3 develops the framework formally: the Complexity and Quality indices built from geometric proxies, the cost and rent height elasticities calibrated from the urban-economics literature, the Monte Carlo machinery, and the joint IRR–DSCR robustness criterion that defines the envelope. Chapter 4 introduces the site, program, and design argument, a constrained parcel in Seattle’s Eastlake corridor developed as a high-rise residential building. Chapter 5 carries out the analysis, scoring the massing schemes against the indices and locating them within the resulting envelope. Chapter 6 steps back to ask what the results imply for practice at the massing stage, for the relationship between architects and developers in speculative production, and for the conditions under which formally ambitious American architecture might again become a defensible category rather than a residual exception. If the earlier chapters establish the model, the last tests whether the discipline can, in fact, stay within its own envelope.

1.5 A Note on What This Thesis Is Not

Before turning to what this thesis sets out to demonstrate, it helps to be clear about what it does not attempt. It is not a critique of the developer, who operates rationally within the constraints capital, cost, and regulation impose. It is not a defense of any particular architectural style; ambition here is a methodological stance, not a formal one. It is not a prediction of project outcomes, since the Monte Carlo machinery is a decision-support tool, not a forecasting device. And it is not a claim that better methodology can substitute for better politics. The conditions under which civic ambition reenters high-budget American architecture are ultimately political, and no feasibility tool will produce them alone. What a feasibility tool can do, and what this one is built to do, is ensure that when those conditions shift, the discipline has the technical vocabulary to respond.

Boom Supersonic has not yet flown a paying passenger above Mach 1. The production aircraft is still in development, the certification path uncertain, and whether commercial supersonic flight returns as a viable category is a question the demonstrator alone cannot settle. This thesis is offered in the same spirit, and with the same limitation. The Safe Design Envelope is a demonstrator. Whether the methodology scales to widespread practice, whether it changes how buildings actually get delivered, is not something this work can resolve on its own. What it can show is that the category is not closed: that the relationship between architectural ambition and financial feasibility, long treated as structurally incompatible, can be articulated within a shared and operational framework, and that the tools to do so are not hypothetical but constructible.

The next chapter establishes the intellectual and empirical ground on which that framework is built. Before the methodology can be developed, the conditions that make it necessary, and the literatures that approach it from adjacent directions, must be set out clearly.

2 Background & Context

Architecture is routinely asked to commit before uncertainty resolves. The decisions that set a project's course, form, scale, and organization, what architects call the *parti*, are made when costs, schedules, markets, and regulatory conditions are still only partly known, yet they are typically justified with analytical tools that presume a certainty the moment does not contain. Adjacent disciplines have spent three decades building methods to describe and manage exactly this uncertainty: real estate finance, computational design, and construction management each have mature, rigorous approaches. Early-stage architectural decision-making in the United States has largely not adopted them, and remains deterministic in structure.

This chapter surveys that work to show how thoroughly uncertainty has been studied in these fields, and why, despite that depth, it has had so little purchase at the schematic stage. The shortfall is not capability. The tools exist, they are mature, and several are already operational in other domains and in peer international markets. The problem is integration. These methods have not been assembled into a form that can operate at the moment of the *parti*, when the geometry is still open to revision and probabilistic financial feedback would do the most good. What follows traces how the strands developed separately, and why their convergence at early design has stayed unresolved.

2.1 Probabilistic Evaluation in Real Estate and Development

Conventional real estate analysis runs on deterministic discounted cash flow (DCF) models that evaluate projects under fixed assumptions about rents, costs, exit values, and discount rates. These models are well understood and deeply embedded in practice, but they handle uncertainty only obliquely, usually by nudging a single discount rate or running one-variable-at-a-time sensitivity tests. The premise is that uncertainty compresses into a handful of parameters yielding a single most-likely outcome, a compression that hides the range, asymmetry, and interaction of possible futures. Two strategies can post the same expected return while carrying entirely different exposure to downside, volatility, or financing failure, differences that decide projects in practice but stay invisible in deterministic appraisal [Hoesli et al., 2006, Loizou and French, 2012]. A spreadsheet is only as honest as its assumptions, and most assumptions are doing more work than they let on.

Leung [2014] develops this more fully, reframing development not as forecasting but as a probabilistic investment problem. His claim is simple and consequential: incorporating uncertainty into the pro forma does not merely change the numbers, it changes how decisions get made. When rent growth, cost escalation, and exit conditions are modeled as stochastic variables rather than fixed values, the objective shifts from maximizing one projected return to managing outcomes across a distribution, and downside risk, tail exposure, and variance become explicit terms of evaluation rather than residue buried in a single net present value. The model stops asking what the answer is and starts asking what the odds are, and how those odds move when the project itself changes. Feasibility analysis becomes naviga-

tion rather than prediction. Hoesli et al. [2006] make the broader methodological case, showing that distributional representations of value and risk outperform point estimates when key inputs are volatile.

This framing reshapes conclusions about timing, phasing, and flexibility. Modeled explicitly, uncertainty turns development from a matter of maximizing expected value under a fixed scenario into one of preserving optionality and weighing trade-offs across plausible futures. The theoretical warrant for treating early flexibility as an asset rather than mere indecision comes from Titman [1985], who shows the option to delay commitment carries measurable value when outcomes are uncertain, and from Dixit and Pindyck [1994], whose treatment of irreversibility supplies the formal reason premature commitment is costly even when present conditions look favorable. Eletter et al. [2024] extend the empirical side, using Monte Carlo simulation to produce interval-based valuations that let decision-makers reason about confidence ranges rather than single points, and Gimpelevich [2011] shows how Monte Carlo return distributions can serve as evaluative signals in high-rise feasibility, classifying design states against return thresholds rather than optimizing toward one predicted outcome. The common thread is that uncertainty does not merely add noise to valuation; it reshapes investment logic.

Scholars caution against mistaking these methods for prediction. Loizou and French [2012] argue Monte Carlo simulation is best understood as a prescriptive framework for structuring judgment under uncertainty, not a forecast of precise outcomes, and locate its limits in data quality, correlated inputs, and the irreducible role of human judgment. Its value depends on how it is embedded in a decision workflow rather than on the sophistication of the simulation itself. That position directly shapes the two-mode structure of Chapter 3, which pairs exploratory simulation with iterative authorship rather than treating Monte Carlo as a generative or optimizing engine.

Design, meanwhile, stays abstracted inside these models. Deterministic and stochastic pro formas alike reduce it to scalars: gross floor area, cost per square foot, revenue assumptions. These can be tuned parametrically, but the geometry that produces them is not itself an analytical variable, so the financial model stays decoupled from the spatial and formal decisions that actually shape risk, cost trajectory, and long-term value. Moorhead et al. [2024] confirm that IRR remains the dominant go/no-go metric in feasibility practice, with hurdle rates typically 12–20% for market-rate multifamily development [Miles et al., 2007], yet the architectural decisions that most directly determine whether those thresholds are met are evaluated separately, if at all.

The decoupling is compounded by the structure of contemporary capital markets. Over two decades, international institutional investment, particularly Australian superannuation funds and Canadian pension vehicles, has reshaped gateway-city development across North America [CBRE Research, 2024, Urban Land Institute and PwC, 2024]. These investors operate at a scale and time horizon that compresses cap rates and inflates land values in high-demand markets, Seattle, Portland, and Vancouver among them. For local developers the result is a land market priced on trophy-asset assumptions while construction costs stay stubbornly local,

a mismatch that narrows feasibility margins and raises the stakes of every parti committed under them. When a parcel trades at a 3.5% implied cap rate before a permit is filed, the gap between a building that pencils and one that does not is measured less in market conditions than in floor-plate efficiency, envelope strategy, and net-to-gross ratio. Parti decisions that read as economically neutral at a 5.5% cap rate become load-bearing at 3.5%. The same capital funds Toronto and Vancouver multifamily, where the resulting buildings are visibly more ambitious in form, which suggests capital structure alone cannot explain the methodological gap between American and peer-market practice. The capital arrives in similar form; what differs is the analytical apparatus that translates it into buildings. This is exactly where a probabilistic, geometry-aware framework has the most to offer.

Work on design value in asset pricing sharpens the disconnect and supplies the second precedent for this thesis. Its foundation is hedonic pricing theory, formalized by Rosen [1974], which decomposes the market price of a differentiated good into implicit prices for its attributes and so makes it possible to isolate what the market pays for specific formal and spatial features. Rong et al. [2020] apply this with unusual specificity to architectural form. Across 3,037 commercial transactions in Manhattan from 2001 to 2017, they find that exterior formal features, diagonality, curvature, setbacks, and podium extrusions, carry statistically and economically significant pricing differentials even after controlling for size, age, class, location, transaction parties, and awarded-architect effects. Diagonal massing commanded a 12.4% premium over rectilinear controls, podium extrusions 9.7%, while setbacks took a 10% discount, with comparable directional effects on net effective rent. Loro et al. [2024] add that daylight availability raises residential prices significantly, independent of location and unit size.

What Rong et al. [2020] establish is that architectural geometry is not economically neutral; it is priced by markets, measurably, with effects on the order of $\pm 10\%$ of transaction value from external form alone. The limit of their framework, and the one this thesis is built to address, is that it works *ex post*: the regression runs on completed, transacted buildings, recovering the implicit prices the market already paid for decisions made years or decades earlier. By the time the coefficient is estimated, the design it describes is fixed. The market already knows form affects value; the design process has not caught up. Rong supplies the empirical warrant for treating form as a financial variable, but not the prospective instrument an architect would need to steer that variable at the moment of the parti, before the building exists and before the differential is baked in.

Together, Leung [2014] and Rong et al. [2020] define the gap. Leung's framework is probabilistic but treats the building as a scalar; Rong's is geometric but treats uncertainty implicitly, as the after-the-fact scatter of transaction prices rather than *ex ante* distribution. Neither lets an architect, at schematic design, ask the question that joins them: *given the stochastic distribution of market conditions the project will meet, and given the geometry of the parti under consideration, what is the probability that this scheme clears its financial thresholds across the full ensemble of plausible futures?* Answering it requires both precedents fused into one instrument, operable at the parti.

2.2 Computational Design and Deterministic Economics

Two decades of computational and parametric methods have widened the capacity for geometric exploration. Parametric workflows generate and test large families of massings against defined performance criteria, while algorithmic tools iterate quickly across constraints such as floor-area efficiency, structural regularity, and envelope performance. Mosey and Deal [2021] show that fundamental geometric variables, building height, width, floor-area distribution, and surface-to-volume ratio, link systematically to construction cost and envelope performance in multifamily housing even before detailed design. That matters here because it establishes massing-stage geometry as a meaningful analytical input rather than a decorative one, the premise on which the proxy framework of Chapter 3 depends. Belniak et al. [2013] similarly find that plan irregularity and perimeter-to-floor-area ratio significantly influence unit construction costs. Together they give empirical grounding to the proposition that parti-stage geometry carries quantifiable cost consequences.

Yet most computational workflows assume stable or predetermined economics. Objectives are framed around spatial feasibility, performance compliance, or geometric optimality, while cost, absorption, and financing are treated as fixed constraints rather than uncertain variables. Optimization runs inside a single assumed economic scenario even when the design process is otherwise exploratory. Geometric exploration and economic assumption remain, for the most part, strangers at the same table.

The limitation shows clearly in early-stage decision-support systems that use stochastic techniques without integrating economic uncertainty. Østergård et al. [2015] propose a stochastic, holistic method for early building design that uses Monte Carlo simulation and sensitivity analysis to explore trade-offs across energy, daylight, and thermal comfort. The approach acknowledges uncertainty in design inputs but keeps it environmental: geometry feeds environmental simulation rather than acting as a driver of economic outcomes, a distinction that looks minor and is not.

More broadly, computational workflows tend to separate geometric decision-making from financial evaluation. Feasibility is assessed after the major moves, the parti, floor-plate configuration, and circulation strategy, are fixed, so economic analysis becomes a verification step rather than a formative input. This mirrors the limitation in probabilistic finance: uncertainty modeled rigorously, form reduced to scalars. Optimizing form is not the same as deciding under uncertainty, and conflating the two has muddled both fields.

2.3 Uncertainty in Construction Delivery

If finance and computational design show the disconnect between uncertainty and form, construction management shows that uncertainty is structural to delivery itself, present whether or not design has acknowledged it. Construction projects are complex, one-off undertakings exposed to technical, organizational, financial, and regulatory risk that upfront planning cannot fully resolve. Reviewing construction

risk practice, Sharma [2020] finds uncertainty across every phase, design coordination, labor, material supply, financing, and regulatory approval, persisting regardless of scale or typology, and frames risk as intrinsic to construction rather than a symptom of poor planning.

The useful distinction in Sharma [2020] is between recognizing uncertainty and operationally treating it. Risk identification and classification are widely adopted; quantitative risk propagation through methods such as Monte Carlo is rarely embedded in early decisions. Uncertainty gets acknowledged descriptively and managed qualitatively, which limits its influence on scope, sequencing, and resource allocation, and the cost is steepest in early stages, where *parti* decisions carry the most downstream weight yet are evaluated on deterministic assumptions.

The temporal asymmetry between when decisions are made and when their consequences surface is captured by the Construction Trilemma, the persistent constraint among scope, quality, and cost in which only two can be actively prioritized at once [Wheeler Kearns Architects, 2014]. Move one and at least one of the others must give. In institutional practice, first cost is often fixed early by budget or financing, leaving scope reduction or quality degradation as the levers for reconciling overruns. Quality has a floor, since life-safety codes impose hard minimums, so impasses recur when cost targets cannot be met by simplification alone. Wheeler Kearns Architects [2014] argue traditional Design–Bid–Build delivery sharpens these impasses, because reliable cost information arrives only after scope and quality are largely committed.




Figure 1: The Construction Trilemma [Wheeler Kearns Architects, 2014].

When misalignment surfaces late, redesign is costly and contentious, with little value recoverable. This tracks the cost-of-change curves of Paulson Jr. [1976], which show the ability to influence cost falling rapidly as a project advances while the cost of modification climbs with growing interdependence among systems. The implication is temporal, not procedural: conventional workflows withhold economic feedback until flexibility has already narrowed. Architectural decisions are

routinely justified under limited risk visibility even though their economic consequences are probabilistic and path-dependent.

Schedule risk illustrates this empirically. Wang and Huang [2009] use Monte Carlo simulation to show that construction schedules carry inherent duration variability: probabilistic variation at the activity level propagates through critical paths to produce a distribution of project durations rather than a single date. Conventional critical-path schedules systematically understate both the likelihood and the magnitude of delay, not through carelessness but because deterministic scheduling cannot represent uncertainty accumulating across interdependent tasks. Burhenne et al. [2010] find the parallel in building performance simulation, where small perturbations in inputs such as occupancy and air-change rates swing simulated energy outcomes, yet deterministic workflows report single values that lend false precision and hide the sensitivity beneath.

Across construction management and building simulation the pattern repeats: uncertainty is acknowledged but structurally suppressed. Schedule models quantify variability without reshaping early design; building simulations expose sensitivity without touching form. Uncertainty peaks at the parti stage, where decisions carry the most downstream influence on cost and value, yet that is exactly where deterministic tools suppress it rather than engage it.

2.4 Set-Based Reasoning as Methodological Precedent

The posture this thesis takes toward uncertainty has a direct precedent in mid-1990s product development, which Chapter 3 develops at length; the relevant point here is brief. Ward et al. [1995] identified set-based concurrent engineering as the distinguishing feature of Toyota's development process, in which designers keep a set of feasible alternatives alive and narrow it only as evidence justifies commitment, rather than committing early to a point design and iterating toward it [Sobek et al., 1999]. Later work has extended these principles across engineering, confirming that set-based strategies outperform point-based ones under uncertainty [Mourtzis et al., 2020]. Ross and Hastings [2005] generalize the logic as tradespace exploration: design as the problem of finding regions of a performance space where multiple solutions stay defensible under varying conditions, rather than locating one optimum in a fixed scenario.

This literature has concentrated on automotive and aerospace manufacturing, but its logic, preserve optionality, explore feasibility regions, and commit only as uncertainty falls, maps cleanly onto architectural design under market uncertainty. Chapter 3 attempts to operationalize that structure at the parti stage, treating the bounded space of plausible massings as the set to be explored before any one configuration is committed.

2.5 The Missing Integration

The preceding sections describe three parallel and largely disconnected trajectories. Real estate finance has shown that probabilistic evaluation yields more meaningful

feasibility signals than deterministic DCF, while still treating geometry as a scalar. Computational design has widened the space of geometric alternatives, but mostly within fixed economic assumptions that neutralize the uncertainty these methods could model. Construction management has established uncertainty as intrinsic to delivery and quantifiable, yet keeps those insights out of early design. Each tradition engages uncertainty seriously; none integrates it with form-giving at the moment the parti is committed.

The thin uptake of uncertainty-aware methods in early American practice reflects not missing capability but a structural mismatch between how the methods are built and how design decisions are actually made. Even sophisticated stochastic design-support systems evaluate geometry against environmental or technical criteria, leaving financial risk outside the design loop. Adjacent construction work makes the gap plain: studies applying Monte Carlo to tolerance and schedule risk show that geometric decisions materially drive risk propagation and downstream cost [Rausch et al., 2019, Li et al., 2016], yet they apply that insight after design is set rather than feeding it back while the parti is still open.

The asset-pricing evidence makes the gap hard to justify. Markets already price architectural decisions, daylight, spatial efficiency, structural regularity, perimeter access, and formal articulation, yet design evaluates those decisions before finance engages them and finance engages them only after design is fixed. The result is a standing misalignment between when architectural choices are made and when their economic consequences are understood. Architecture commits early, finance arrives late, and the project absorbs the difference.

The misalignment is not evenly distributed. The peer markets named in the introduction, Toronto, Melbourne, Amsterdam, Istanbul, and Singapore, keep producing mid- and high-rise residential at a scale and level of formal resolution that shows the integration is operationally achievable. This is not a solved condition. In many of those contexts the same forces that widen the feasible set also standardize its expression, yielding a high volume of well-executed buildings within a relatively narrow architectural language shaped by delivery efficiency and capital flow. Their relevance is not as models to copy but as existence proofs: architectural ambition at scale, under speculative capital and real cost conditions, is not a closed problem.

Despite extensive work on uncertainty in adjacent fields, American architectural practice lacks a formal way to decide under uncertainty at the parti stage. Early choices are justified by precedent, intuition, or deterministic metrics even when their economic consequences are probabilistic and path-dependent. The gap is not methodological but epistemic: the discipline lacks a consistent way to operate when outcomes cannot be predicted, only bounded. Existing methods either model uncertainty without geometry or generate geometry without uncertainty.

This thesis addresses that missing integration by treating architectural form not as a static input to financial models but as an active variable within probabilistic evaluation, authored, measured, and tested against stochastic market conditions at the parti stage, before commitment narrows the field. The next chapter makes that integration operational.

3 Methodology

At the massing stage an architect commits most of a building's cost long before anyone prices it. The parti fixes footprint, profile, core, and envelope, and with them the structural and façade burdens the pro forma will only later discover. Yet the two decisions are usually made in sequence and by different hands: form is resolved first, feasibility is checked after, and whatever conversation occurs between them tends to arrive as a series of value-engineering retreats. The premise of this thesis is that the design decision and the investment decision are the same decision, and that they are most usefully made together, while the geometry is still soft.

The standard response to market uncertainty is to optimize toward a single predicted future: one pro forma, one set of assumptions, one answer that is either right or, more often, brittle. This framework declines that bet. Rather than search for the one scheme that pencils under one forecast, it maps the region of design states, combinations of architectural complexity (C) and asset quality (Q), that stay financially defensible across an ensemble of plausible futures. That region is the Safe Design Envelope (SDE): not a point to aim at but a space to stay inside, the set of formal positions the market can be expected to forgive. The move is from point-based to set-based reasoning about design under uncertainty, committing late and only as evidence accumulates [Sobek et al., 1999].

The framework operates through two complementary modes. *Mode A* performs a stochastic sweep across a bounded space of abstract design states defined by ranges of C and Q , evaluated against stochastic realizations of the market environment Ω ; its output is the SDE itself. *Mode B* reintroduces architectural authorship as a focused, iterative act within that envelope: discrete configurations are designed, encoded through geometric proxies, and evaluated probabilistically, with diagnostics such as expected value, variance, and downside exposure informing successive adjustments. The result is a feedback loop between architectural decision-making and probabilistic financial evaluation.

The methodology that follows operationalizes this logic through a framework requiring three capabilities. First, architectural decisions at the massing stage must be encoded into quantitative signals that can enter a financial model; this motivates the Complexity and Quality indices (§3.3–§3.5). Second, design must be evaluated against a range of market conditions rather than single-point assumptions; this motivates the Monte Carlo simulation (§3.8). Third, the framework must distinguish robust configurations from vulnerable ones in terms that connect to development decision-making; this motivates the joint IRR-DSCR criterion and the SDE (§3.9). The overall structure follows the nested modeling hierarchy Kim [2020] formalize for construction-project risk management (a generation model, a simulation model, an optimization model, and a risk-management model), which Lee and Cha [2019] and Rachmawati and Kim [2023] implement for financial risk factors in Korean apartment development. The present framework transposes that structure to design-state variables, coupling architectural encoding to probabilistic financial evaluation without the broader systems-analysis apparatus, which addresses concerns outside the scope of architectural decision support.

Model Inputs and Outputs

The mathematical structure of the model is not a precise predictive representation of development outcomes. It is a deliberately simplified, parametric abstraction that preserves the dominant relationships between architectural form, construction cost, and revenue generation while enabling tractable exploration of uncertainty and design sensitivity. Where established formulations exist in the finance, urban-economics, or construction-management literature, they are adopted directly and cited. Where adaptation is required to admit geometric design variables that prior models do not accommodate, the modification is stated explicitly and its rationale documented. Where constructs are introduced without direct precedent, their necessity is justified by the absence of existing instruments that couple design encoding to stochastic financial evaluation.

Three categories of quantity run through the entire model and should not be confused. *Measured geometry* comprises values a Rhino model computes deterministically: every Complexity sub-metric and the geometric Quality sub-metrics are extracted, not estimated, and carry no probability distribution. *Interpreted design judgment* comprises the three strategic scores, authored by the designer to register intent that geometry cannot encode; these are also deterministic for a given scheme. *Financial and market assumptions* comprise the stochastic vector Ω , the only layer of the model that is sampled rather than computed. The distinction conditions everything that follows: design metrics are reported with analytic bounds, while the market variables are the only quantities for which a mean, standard deviation, and distributional family are specified. The provenance and bounds of every sub-metric are recorded in Table 3 (§3.5); the distribution of every market variable is recorded in Table 10 (§3.8).

3.1 Financial Model

Each simulation run assembles a project scenario from three input groups.

A **market condition** is sampled from the stochastic market-state vector Ω , comprising variables outside the designer’s control: construction cost escalation, rent growth, absorption rates, financing terms, exit capitalization rates, and schedule uncertainty. These variables are not statistically independent, and ignoring their correlation can generate economically incoherent scenarios [Touran and Wiser, 1992].

An **architectural design state** is specified through the pair (C, Q) . In Mode A these values are sampled directly as abstract intensities across a bounded grid; in Mode B they are computed from designed geometric configurations through the extraction pipeline of §3.5.

A **fixed project context** — site characteristics anchor all simulations to a common baseline so that observed differences in outcome are attributable to design decisions and market uncertainty rather than to changes in project definition. A second control operates on the market itself: when design alternatives are compared, the framework holds not only the project context but the sampled futures constant. Every scheme in a Mode B comparison is evaluated against the same N




Figure 2: Methodological flowchart

market draws from Ω , a paired construction detailed in §3.11. Two schemes therefore face identical futures, and the difference between their outcomes is attributable to the design move alone rather than to sampling noise — the discipline a controlled comparison requires.

3.2 3D Experiment Space

The simulation begins within a Rhino 7 environment designed to decompose architectural form into an array of performance metrics. By tying the geometry to Grasshopper, every formal move and parametric exploration can be instantly quantified.

Figure 3 shows the basic extraction workflow. Rhino geometry enters the GHPython script as the measurable side of the experiment, while a small set of qualitative inputs records design judgments that geometry alone cannot carry. The distinction matters. A floor plate can tell the script its area, perimeter, and relationship to the core; it cannot tell the script whether the scheme is materially ambitious, environmentally deliberate, or attempting a higher market position. Those judgments remain authored inputs.



Figure 3: Grasshopper workflow

The workflow is organized around two streams. The first is geometric: surfaces, Breps, curves, and areas that Rhino can measure directly. The second is qualitative: scalar inputs representing design ambitions visible in an architectural proposal but not automatically recoverable from its mass. Both streams pass through the GHPython script and emerge as composite indicators. In practical terms, this creates a feedback loop between model and pro forma: a change to the tower’s shape, core position, or perimeter condition is no longer a purely formal adjustment, but a change in the project’s measurable design state. The specific Rhino inputs used in this extraction are defined in Table 1.

Site topography Topography enters the model as a scalar complexity index $T \in [0, 1]$, computed once from the project surface in Rhino/Grasshopper. The surface is discretized into a uniform grid; average slope and elevation variability are extracted (the latter normalized by the site’s horizontal bounding-box diagonal for scale consistency) and combined into a single bounded value through fixed weights. The result is a lightweight proxy for grading difficulty, foundation complexity, and

Table 1: Rhino geometric inputs and modeling definitions.

Name	Description in Rhino terms
Floorplates	Planar Breps or surfaces representing each horizontal floor plate, typically generated by sectioning the massing at each level.
Leasable	Planar Breps representing each floor’s leasable area, drawn around the units and excluding core and back-of-house.
Core	Closed Brep volume representing the vertical service, circulation, and structural core.
Skin	Exterior Brep or joined surface defining the outside envelope of the massing, excluding floor plates and internal geometry.
Core footprint	Planar Brep or closed planar curve representing the core footprint at a typical floor level.
Site area	Planar closed boundary curve or Brep defining the site footprint after property-line setbacks.

site logistics, in keeping with established geomorphometric practice that treats basic statistical descriptors of terrain as effective measures of surface complexity [Grohmann et al., 2011]. As an exogenous input, T is fixed across simulations and enters the cost model as a multiplicative site-condition factor $(1 + \kappa_T T)$, where $\kappa_T = 0.15$ represents a stylized moderate impact, on the order of 5–20% for typical urban sites [Huang, 2021].

Performance Diagnostics

Each assembled project state is passed to a deterministic financial engine that returns four performance metrics: Net Operating Income (NOI), Yield-on-Cost (YOC), Internal Rate of Return (IRR), and Debt Service Coverage Ratio (DSCR). NOI is the foundational cash-flow figure from which downstream metrics are derived. YOC relates stabilized NOI to total project cost, producing a development spread against the exit capitalization rate sampled from Ω ; because higher complexity raises total cost, YOC provides a direct channel through which C influences viability independent of revenue assumptions. IRR is the primary robustness criterion, reflecting its standing as the go/no-go metric in speculative development [Moorhead et al., 2024].

DSCR is the complementary diagnostic that surfaces financing-side failure modes IRR alone may miss: a configuration can exhibit attractive expected returns while becoming debt-service insolvent when cost escalation and rent underperformance co-occur. The DSCR threshold (τ) used in this study ($\tau_{\text{DSCR}} = 1.25\times$) is not arbitrary. It reflects institutional underwriting practice: Fannie Mae’s conventional multifamily term sheet specifies a minimum DSCR of $1.25\times$ [Fannie Mae, 2026], while Freddie Mac’s Small Balance program requires amortizing coverage ratios

of $1.20\times$ to $1.40\times$ by market tier [Freddie Mac, 2026]. A project whose DSCR distribution breaches these thresholds under stress is not merely underperforming; it is non-bankable under standard agency financing, regardless of its equity-return profile.

3.3 Architectural Complexity (C)

Complexity and Quality are the two composite indices through which design enters the financial model. Together they are assembled from ten geometric and qualitative proxies: this thesis’s way of abstracting architecture and design into quantifiable signals, so that formal decisions can enter a stochastic pro forma at all. The abstraction rests on two strands of work. The first establishes that design is a measurable, market-priced attribute rather than an intangible: hedonic theory reads a building as a bundle of attributes the market prices implicitly [Rosen, 1974], and computational valuation has since put numbers on specific external form features and their separable effect on price [Rong et al., 2020]. The second establishes that a development venture is better read as a distribution than as a single figure: probabilistic, simulation-based evaluation has been advanced as a corrective to deterministic discounted-cash-flow analysis, which compresses an uncertain venture into one number and hides the very risk a developer is paid to manage [Leung, 2014, Geltner and de Neufville, 2018]. This thesis sits where the two strands meet, running quantified design forward through a probabilistic pro forma so that a massing decision and an investment decision become the same decision.

Complexity is the cost-facing half of that abstraction. It captures how geometric and organizational decisions amplify construction effort, coordination demand, and exposure to cost uncertainty, and it draws on the proxies that govern cost rather than value. For any parti’s specific design configuration X_i , a pre-height geometric score $C_G(X_i)$ aggregates those proxies into a single dimensionless reading; §3.6 defines and weights each one.

The geometric proxies isolate formal properties of the massing independent of absolute scale, so two buildings with identical floor-plate geometry, equivalent taper, and similarly organized cores register the same C_G regardless of height, and an identification check confirms the proxies are largely height-independent. That invariance is deliberate, since formal complexity and the scale-dependent economics of building tall are distinct mechanisms. The height consequence is added afterward rather than mixed in: §3.7 exposes C_G through the slenderness aspect ratio to the cost-facing index C , so height is internalized through the geometry it actually changes rather than applied as a separate elasticity. This keeps the model consistent with the skyscraper-economics literature, where the marginal cost of height is real but enters unevenly across building systems rather than as a single proportional surcharge [Ahlfeldt and Barr, 2022, Ahlfeldt and McMillen, 2018]. Within the simulation, C modulates both the level and the dispersion of construction-cost outcomes: higher complexity shifts cost distributions upward and widens their spread, amplifying vulnerability to cost volatility rather than predicting cost directly.

3.4 Asset Quality (Q)

Quality is the value-facing half of the same abstraction, built from the remaining proxies. Where Complexity reads how design amplifies cost, Quality reads how it generates value, across two registers. A geometric register, read directly from the massing, establishes baseline spatial viability through the properties markets reliably capitalize, among them programmatic yield, spatial efficiency, and daylight access. A qualitative register, designer-assigned, carries the positioning intent the geometry cannot, among them sustainability ambition, programmatic differentiation, and material and tectonic quality. Both are financially consequential, and holding only one in view leaves the building's market position half-read.

The split tracks what the market can actually read. Legible, geometrically expressible attributes are capitalized into value while less perceptible ones are not [Nase et al., 2016], so the framework routes the former through the geometric register and treats the rest as strategic. That the qualitative register is authored rather than extracted is a defended choice, not a concession: developers routinely recognize sustainability and experiential value before the market measures it cleanly, the condition under which an authored input is the right instrument [van Bakergem, 2020].

The two registers do not enter the pro forma as a single number. Because height bears differently on rentable area and on rent rate, §3.7 routes the yield and efficiency readings into an area-facing transfer score Q_A and the daylight and qualitative readings into a rate-facing transfer score Q_R , and the revenue model reads those two channels separately. Their effect is scaled by the market-level revenue premiums documented in §4.5 rather than by the quality scores themselves, the two operating serially: a quality position improves competitiveness and compresses downside exposure, while the market sets what that position is worth. The composite index Q , the harmonic mean of Q_A and Q_R , is a reporting measure and the quality ordinate of the design envelope. Unlike C , which enters the cost model directly, Q describes rather than drives. The asymmetry is intentional: cost answers to one composited burden, while value reaches the model through the two channels height pulls apart.

3.5 The Proxies

This study treats the move from geometry to metric as a measurement claim with explicit limits, not as a proxy for architectural value. Yoo et al. [2025] demonstrate that architectural characteristics can be operationalized into reproducible computational metrics without task-specific retraining, but caution that such metrics quantify formal properties, not design quality or economic value. C and Q inherit this constraint by construction: they encode geometric and organizational structure that the construction-economics and hedonic literatures link to cost and revenue mechanisms; they do not encode merit. The financial model, not the index, carries the economic interpretation, and the indices are admissible only insofar as their inferred structure survives the robustness testing of §3.10.

The indices C and Q are computed programmatically via a GHPython component in Grasshopper that queries the Rhino model at each evaluation step, so that as configurations are revised in Mode B, scores are recomputed consistently from the same geometric source. Four inputs are required: planar surfaces as floor plates produced by horizontal sectioning of the massing, an open exterior skin surface, a planar surface representing the service-core footprint at a representative floor level, and a closed volume representing the volumetric extent of the service core through the building section. The full proxy set is summarized in Table 2; provenance and analytic bounds for every sub-metric are recorded in Table 3 at the end of this section.

Table 2: Proxy variables constituting the Complexity and Quality indices.

Variable	Design interpretation
<i>Complexity variables (C)</i>	
C_{env}	Envelope Burden: envelope area relative to enclosed volume; captures envelope inflation and facade burden.
C_{stack}	Stack Variation: variation in floor-plate area across levels; captures taper, setbacks, and podium transitions.
C_{edge}	Edge Variation: variation in floor perimeter across levels; captures boundary mutation through the section.
C_{core}	Core Offset: displacement of the service core from the floor centroid; captures structural asymmetry.
<i>Quality variables (Q): geometric, Q_{geo}</i>	
Q_{yield}	Yield: efficiency of footprint occupation independent of total height.
Q_{eff}	Net-to-Gross Efficiency: spatial efficiency of the floor plate; proportion of constructed area that remains leasable.
Q_{sun}	Solar Exposure: perimeter exposure to daylight-favorable orientations.
<i>Quality variables (Q): qualitative, Q_{qual}</i>	
Q_{green}	Green Score: designer-assigned score for green space, high-performance systems, and environmental design ambition.
Q_{innov}	Innovation Score: designer-assigned score for advanced construction materials, systems, or methods.
Q_{mat}	Material Score: designer-assigned score for tectonic ambition and the perceived quality of the assembly.

Proxies of Architectural Complexity

Four geometric sub-metrics constitute the Complexity Index. Each addresses a distinct aspect of how architectural decision-making at the massing stage accumulates construction effort, coordination risk, and delivery uncertainty.

Envelope Burden (C_{env}) measures the surface-to-volume condition of the massing. Derived from the Surface-to-Volume Ratio S/V , it relates exterior envelope area to enclosed building volume, capturing the extent to which formal decisions inflate the building envelope beyond what is required for a given program. D’Amico and Pomponi identify the relationship between external surface area and inter-

nal volume as a geometric basis for evaluating building-form compactness, while noting that raw S/V is scale-dependent and therefore requires normalization or bounded comparison when used across alternatives [D’Amico and Pomponi, 2019]. Their aspect-ratio framework further clarifies that envelope burden is not just an output of surface area, but is produced by measurable geometric tendencies such as plan elongation, vertical slenderness, and departure from compact reference forms. In this thesis, C_{env} is retained as a raw geometric proxy at the extraction stage and normalized later with the other complexity variables before aggregation into C .

$$C_{\text{env}} = \frac{A_{\text{skin}}}{V}. \quad (1)$$

Compact, prismatic towers minimize C_{env} by enclosing volume efficiently. Massings with elongated floor plates, high slenderness, stepped profiles, articulated facades, cantilevers, or distributed volumes increase envelope area without proportionally increasing enclosed volume, raising the envelope-burden score. This interpretation follows D’Amico and Pomponi’s argument that reducing envelope surface area for a given amount of internal volume improves form efficiency and can reduce the material and energy burdens associated with the building envelope [D’Amico and Pomponi, 2019]. Accordingly, C_{env} serves as a general proxy for enclosure-driven cost and resource escalation.

Stack Variation (C_{stack}) measures the vertical profile of the parti, and captures variation in gross floor-plate area across levels, translating setbacks, stepbacks, podium transitions, cantilevered volumes, and other departures from a prismatic stack into a single relative-dispersion measure. A tower that maintains a consistent floor plate from ground to roof reduces geometric discontinuity, supports repetition in formwork and facade coordination, and limits the number of structural and mechanical transitions that must be resolved through the section. Conversely, variation in floor-plate area introduces discontinuity between levels, increasing the coordination burden associated with structure, envelope, services, and construction sequencing.

Letting F denote the total number of floors indexed by $i = 1, \dots, F$, with floor areas $A_{\text{floor},i}$, C_{stack} is calculated as the coefficient of variation of floor-plate area:

$$C_{\text{stack}} = \frac{\text{std}\{A_{\text{floor},i}\}_{i=1}^F}{\text{mean}\{A_{\text{floor},i}\}_{i=1}^F}. \quad (2)$$

A prismatic stack produces $C_{\text{stack}} = 0$; the index rises with the severity and frequency of departures from floor-plate uniformity. Because the coefficient of variation divides dispersion by the mean, C_{stack} measures proportional variation rather than absolute area change, allowing equivalent relative taper or stepping to register consistently across differently sized alternatives [National Institute of Standards and Technology, 2017]. This interpretation is consistent with construction-cost research showing that building shape and dimensional variation affect cost-relevant quantities such as walls, foundations, usable area, and unit construction cost [Belniak et al., 2013].

Edge Variation (C_{edge}) measures perimeter variability across the massing. It captures fluctuations in exterior floor-plate perimeter across levels, translating chamfers, rotations, re-entrant notches, changing balcony edges, and other boundary transformations into a single relative-dispersion measure. Where C_{stack} measures changes in floor-plate area, C_{edge} asks whether the boundary of the floor plate is itself changing. A building whose plates shrink uniformly may produce perimeter variation that largely mirrors area variation. A building whose boundary actively transforms introduces additional edge-condition changes that must be resolved through facade geometry, panelization, detailing, and construction sequencing.

Again with F denoting the total number of floors indexed by $i = 1, \dots, F$, let exterior floor-plate perimeters be denoted by P_i . C_{edge} is calculated as the coefficient of variation of floor-plate perimeter:

$$C_{\text{edge}} = \frac{\text{std}\{P_i\}_{i=1}^F}{\text{mean}\{P_i\}_{i=1}^F}. \quad (3)$$

A prismatic tower with a consistent floor-plate boundary produces $C_{\text{edge}} = 0$; the index rises with the severity and frequency of perimeter changes through the section. Because the coefficient of variation divides dispersion by the mean, C_{edge} measures proportional perimeter variability rather than absolute perimeter length, allowing differently sized alternatives to be compared on the basis of relative boundary change [National Institute of Standards and Technology, 2017]. This interpretation is consistent with facade cost research showing that facade-system cost is shaped by factors such as facade area, building height, and the complexity of the building body [Leśniak and Górka, 2021].

Core Offset (C_{core}) measures the eccentricity of the service core relative to the floor-plate centroid. Derived from the Awida’s Core Eccentricity study (2024), it captures the organizational and structural consequences of moving the core away from the geometric center of the plan. A centered core helps regularize structural bays, distribute gravity loads more symmetrically, and align the plan more closely with the centric-core condition commonly assumed in early tall-building studies. As the core migrates toward the perimeter, the distance between the building’s geometric center and its principal service-and-lateral-resistance zone increases, creating a plan imbalance that may affect wind drift, dynamic behavior, and the need for corrective structural coordination.

Let F denote the total number of floors indexed by $i = 1, \dots, F$, with floor-plate centroids $\mathbf{c}_{\text{floor},i}$, service-core centroid $\mathbf{c}_{\text{core}}$, and half-diagonal floor-plate bounding-box length k_i . C_{core} is calculated as the average relative distance between the floor centroid and the service-core centroid:

$$C_{\text{core}} = \frac{1}{F} \sum_{i=1}^F \frac{\|\mathbf{c}_{\text{floor},i} - \mathbf{c}_{\text{core}}\|}{k_i}. \quad (4)$$

A centered core produces $C_{\text{core}} = 0$; the index rises monotonically as the core migrates outward relative to the floor plate. Dividing centroid displacement by k_i allows core offset to be compared across differently sized floor plates, since the score measures relative displacement rather than absolute distance. Awida’s study of a 55-floor concrete tall building compares centric core placement with cores offset by one sixth and one third of the building width, finding that the one-third-width eccentric case produced substantially higher wind drift, longer fundamental periods, and higher peak acceleration, eventually requiring secondary cores to restore acceptable behavior [Awida, 2024]. Accordingly, C_{core} serves as a massing-stage proxy for core eccentricity and the structural coordination risk associated with plan imbalance.

Proxies of Asset Quality

Yield (Q_{yield}) measures the average programmatic yield of the massing relative to the fixed site footprint. Derived from the industry standard Floor Area Ratio (FAR), it evaluates how much net leasable area each floor delivers per unit of land without rewarding vertical scale. FAR is useful as a zoning and development control metric because it relates total building floor area to parcel area, but it conflates scale with quality by aggregating floor area across all stories [Seattle City Council, 2019]. A taller building can therefore receive a higher FAR even when its individual plates are inefficient. Q_{yield} instead isolates the geometric efficiency of the configuration itself:

$$Q_{\text{yield}} = \frac{1}{F} \sum_{i=1}^F \frac{A_{\text{net},i}}{A_{\text{site}}}, \quad (5)$$

where $A_{\text{net},i}$ is the net leasable area of floor i , A_{site} is the fixed site area, and F is the total number of floors. Normalizing each plate individually against site area and then averaging rewards footprint strategies that maximize usable space per unit of land without treating height itself as a quality signal. This decoupling is empirically warranted. Dong [2023] finds that the relationship between multifamily density and home prices in the Portland metropolitan area is not mechanically monotonic: medium-density multifamily homes sell at lower prices than other housing types, the cost-saving effect of higher density development is weak or negative, and dwelling size is the primary determinant of multifamily home prices. Development-feasibility practice likewise evaluates project viability through both construction costs and the future income stream, meaning that yield should be interpreted through usable, revenue producing area rather than through total bulk alone [Local Housing Solutions, n.d.]. Q_{yield} is constructed to avoid importing total vertical accumulation into the quality signal.

Spatial Efficiency (Q_{eff}) measures the proportion of constructed floor area that remains directly leasable. Derived from the common measure of Net-to-Gross Efficiency, it captures the financial and spatial consequence of converting gross floor

area into revenue producing area. Every square foot lost to cores, corridors, shafts, service zones, and unusable residual space still costs money to build but does not directly produce rent. As a result, Q_{eff} functions as a practical efficiency proxy linking plan organization to feasibility.

$$Q_{\text{eff}} = \frac{\sum_{i=1}^F A_{\text{lease},i}}{\sum_{i=1}^F A_{\text{floor},i}}. \quad (6)$$

Here, $A_{\text{lease},i}$ is the measured leasable area on floor i , calculated from the summed unit-area or leasable-area geometry rather than inferred by subtracting the core from the floor plate. This distinction matters because the non-leasable remainder may include more than the core alone, such as corridors, shafts, service areas, unusable residual spaces, and floor-specific irregularities. In cases where direct leasable geometry is unavailable, $A_{\text{floor},i} - A_{\text{core},i}$ may be used only as a fallback approximation, not as the primary measurement.

This formulation treats spatial efficiency as a measurable development condition rather than as a stylistic preference. Real estate area measurement practice distinguishes usable, rentable, gross, and common area relationships because the conversion between usable area and rentable area affects leasing comparison, valuation, benchmarking, and operating economics [NAIOP, n.d., Building Owners and Managers Association International, n.d.]. HUD multifamily guidance similarly distinguishes net residential rentable area, gross floor area, and area not attributable to dwelling use, confirming that the relationship between dwelling-use area and total building area matters in multifamily cost and underwriting contexts [U.S. Department of Housing and Urban Development, 2023]. Recent building performance research further supports the importance of spatial efficiency: Weber et al. [2024] find that effective space use can outperform envelope upgrades for operational carbon reduction in substantial portions of surveyed residential buildings across Zurich, New York, and Singapore. The Q_{eff} proxy therefore carries both financial and environmental significance.

Solar Exposure (Q_{sun}) measures the solar orientation quality of a scheme's exterior perimeter. Derived from the Passive Daylight Index, it operates as a length-weighted orientation exposure proxy rather than a full daylight simulation. Because the unit catalog places residential units within conventional daylight depth, a binary depth-band test saturates and cannot differentiate schemes. Q_{sun} therefore evaluates how much of the exterior perimeter is oriented toward more daylight rich directions.

For each exterior perimeter segment i , the outward normal is compared to true north, yielding a continuous orientation score $s_i \in [0, 1]$ from the angular deviation $\theta_i \in [0, \pi]$:

$$Q_{\text{sun}} = \frac{\sum_{i=1}^n L_i s_i}{\sum_{i=1}^n L_i}, \quad s_i = \frac{\theta_i}{\pi}, \quad (7)$$

where L_i is the segment length and n is the number of exterior perimeter segments. North-facing frontage scores lowest, east- and west-facing frontage intermediate,

and south-facing frontage highest. The angular formulation avoids the cancellation that a linear cosine expression produces on a closed perimeter, where opposing normals force the length-weighted average toward a neutral value regardless of orientation. Given the project scope, a full climate-based daylight simulation is inappropriate; however, daylight availability remains a recognized residential performance concern, and simulation-based metrics such as spatial daylight autonomy are typically used when a more detailed daylight analysis is warranted [Hong et al., 2024, Solemma, 2024]. Q_{sun} therefore provides a compact massing-stage proxy for relative solar-orientation exposure.

Strategic Scores The qualitative Asset Quality variables Q_{green} , Q_{innov} , and Q_{mat} are assigned on a 1–10 ordinal scale using fixed rubric anchors. Intermediate values may be used when a scheme falls between two anchors. These scores are not empirical measurements; they are structured judgments intended to make strategic design intent explicit and repeatable within the experiment.

- **Score of 1: Minimal.** The scheme shows little evidence of the strategic category beyond a conventional baseline.
- **Score of 3: Limited.** The scheme includes isolated or secondary gestures, but they do not materially shape the project.
- **Score of 5: Competent.** The scheme demonstrates a clear, credible strategy, but one that remains conventional or expected for the project type.
- **Score of 7: Strong.** The strategy is legible and materially affects the architectural or market identity of the proposal.
- **Score of 10: Project-defining.** The strategy is central to the design concept and would be difficult to remove without changing the project’s identity.

Applied to the three strategic sub-metrics, these anchors can be interpreted as follows:

- **Green Score (Q_{green}).** A score of 1 indicates code-minimum environmental performance with little visible sustainability strategy. A score of 5 might include planted terraces, shading, operable windows, stormwater features, or a coherent energy concept [LEED or WELL certification, etc.]. A score of 10 would require environmental strategy to define the project, for example through substantial biophilic integration, low-carbon structural or cladding systems, meaningful landscape area, water strategy, or visible resilience features.
- **Innovation Score (Q_{innov}).** A score of 1 indicates a standard stacked-unit product with little differentiation. A score of 5 might include a stronger

amenity concept, improved unit variety, a clearer mixed-use base, or a recognizable formal identity. A score of 10 would require the spatial or programmatic concept to drive the project’s market position, for example through a distinctive residential typology, unusual shared-space strategy, novel unit organization, or a massing concept that produces recognizable value rather than novelty alone.

- **Material Score** (Q_{mat}). A score of 1 indicates a neutral or thinly resolved envelope with limited material specificity. A score of 5 might include durable materials, greater facade depth, clearer balcony or shading details, or a coherent base-to-tower expression. A score of 10 would require material and tectonic logic to become central to the proposal, for example through expressive structure, high-quality cladding, mass timber expression, facade depth, or an assembly strategy that communicates durability, craft, and specificity.

3.6 Normalization and Composite Scoring

Every Complexity proxy is an empirical geometric measurement of the parti, and every such reading begins life nonnegative: $c_j(X_i) \in \mathbb{R}_{\geq 0}$. A building cannot post a negative envelope burden, a negative floor-plate variation, a negative perimeter fluctuation, or a core somehow more centered than centered. The raw values C_{env} , C_{stack} , C_{edge} , and C_{core} are normalized after extraction and before they are aggregated into a pre-height geometric Complexity score. Each reading is drawn into a bounded, monotone design-intensity score by a reference-saturation function, a single-reference value function in the multi-attribute tradition [Keeney and Raiffa, 1976] and a familiar bounded normalization in composite-indicator practice [OECD and European Commission, Joint Research Centre, 2008]:

$$\hat{c}_j(X_i) = \frac{c_j(X_i)}{c_j(X_i) + r_j}, \quad r_j > 0, \quad \hat{c}_j(X_i) \in [0, 1). \quad (8)$$

where r_j sets the reference intensity, the point at which a proxy reaches a moderate and materially meaningful expression within the design space of this experiment. At $c_j = r_j$ the normalized score reads exactly 0.5. Below the reference the gesture is still latent; above it the score climbs toward 1 and, like a tower chasing its own asymptote, never quite tops out. The form spares us from declaring a universal ceiling for burdens that have no natural one, among them envelope burden, stack variation, edge variation, and core offset.

The normalized complexity variables $\hat{C}_{\text{env}}$, $\hat{C}_{\text{stack}}$, $\hat{C}_{\text{edge}}$, and $\hat{C}_{\text{core}}$ are then aggregated under a softened rank-decay weighting. Ordinal surrogate weights such as rank-sum, rank-reciprocal, and rank-order centroid are well established for cases where the criteria can be ranked with confidence but precise weights cannot, and they prove surprisingly hard to beat when the aim is choosing the better alternative rather than pricing it to the penny [Stillwell et al., 1981, Barron and Barrett, 1996]. The four proxies are ordered by their expected bearing on early feasibility. Stack Variation leads, since floor-plate variation is what unsettles repetition, stacking, and

the quiet economies of a building that wants to be the same all the way up. Core Offset follows, because an eccentric core carries structural and organizational consequences and has become a deliberate, increasingly common move in tall buildings rather than an accident [Oldfield and Doherty, 2019]. Edge Variation sits third as a secondary reading of perimeter mutation, and Envelope Burden sits last, still in the room but not yet running it at the massing stage.

To soften the cliff of binary decay without resorting to arbitrary decimals, the ranking is mapped onto a square-root sequence ($\sqrt{4} : \sqrt{3} : \sqrt{2} : \sqrt{1}$), a concave middle path between the linear rank-sum and the steeper rank-reciprocal schemes [Stillwell et al., 1981]. The pre-height geometric Complexity score is:

$$C_G(X_i) = \frac{\sqrt{4} \hat{C}_{\text{stack}}(X_i) + \dots + \sqrt{1} \hat{C}_{\text{env}}(X_i)}{\sqrt{4} + \dots + \sqrt{1}} \longrightarrow \frac{\sum_{r=1}^4 \sqrt{5-r} \hat{c}_r(X_i)}{\sum_{r=1}^4 \sqrt{5-r}} \quad (9)$$

where r is the priority rank of the normalized proxy $\hat{c}_r$. The sequence resolves to proportional weights of 32.5% for $\hat{C}_{\text{stack}}$ ($r=1$), 28.2% for $\hat{C}_{\text{core}}$ ($r=2$), 23.0% for $\hat{C}_{\text{edge}}$ ($r=3$), and 16.3% for $\hat{C}_{\text{env}}$ ($r=4$). C_G reads the parti's formal and organizational complexity before height has its say. It is not yet the cost-facing transfer score; that office falls to C_T in §3.7.

Quality is treated as a *synergistic* condition, with a clean separation between geometric spatial viability and qualitative strategic differentiation. Before height enters, the three geometric quality proxies, programmatic yield, spatial efficiency, and solar exposure, are taken as jointly necessary, since a shortfall in any one of them quietly sinks leasability no matter how the others perform. Their pre-height spatial reading is aggregated multiplicatively, through a geometric mean:

$$Q_{\text{geo}}(X_i) = \sqrt[3]{\hat{Q}_{\text{yield}}(X_i) \cdot \hat{Q}_{\text{eff}}(X_i) \cdot \hat{Q}_{\text{sun}}(X_i)}. \quad (10)$$

This non-compensatory form makes Q_{geo} the quality-side counterpart to C_G . A weakness in any single geometric dimension pulls the whole score down rather than being papered over by strength elsewhere, which sits with the long-standing reminder that the choice of aggregation is never neutral and that a geometric mean refuses the easy substitutions an arithmetic average permits [OECD and European Commission, Joint Research Centre, 2008, European Commission, Joint Research Centre, n.d.]. The United Nations Development Programme made precisely this switch in the 2010 Human Development Index, trading the arithmetic mean for the geometric one so that a country could no longer buy its way out of one poor dimension with a strong showing in another [United Nations Development Programme, 2010]. Since height bears asymmetrically on rentable area and on rent rate, the operative aggregation happens channel by channel in §3.7, where yield and efficiency form the area-facing transfer score Q_A and solar exposure is carried into the rate-facing transfer score Q_R . Equation (10) stays on as the pre-height spatial reading kept for reporting.

The qualitative component Q_{qual} enhances spatial viability rather than defining it, and is aggregated additively as an equal-weighted mean:

$$Q_{\text{qual}}(X_i) = \frac{\hat{Q}_{\text{green}}(X_i) + \hat{Q}_{\text{innov}}(X_i) + \hat{Q}_{\text{mat}}(X_i)}{3}. \quad (11)$$

Equal weighting is adopted in the absence of any evidence that would justify favoring one of the three strategic legs over the others, as they are deliberately abstracted. The reporting Q is left unfinished here on purpose. Height reshapes the area-facing and rate-facing components in §3.7, so the normalized quantities above stand as the pre-height inputs to that transfer sequence.

Table 3: Proxy provenance, normalization, and structural weights.

Metric	Provenance	Normalization Logic	Weight / Aggregation Rule
<i>Architectural Complexity</i> (C)			Square-Root Rank Decay (Eq. 9)
$\hat{C}_{\text{stack}}$	Measured / Adapted	Ref. saturation ^a	32.5% (Rank 1)
$\hat{C}_{\text{core}}$	Measured / Authored	Ref. saturation ^a	28.2% (Rank 2)
$\hat{C}_{\text{edge}}$	Measured / Adapted	Ref. saturation ^a	23.0% (Rank 3)
$\hat{C}_{\text{env}}$	Measured / Adapted	Ref. saturation ^a	16.3% (Rank 4)
<i>Asset Quality</i> (Q)			Harmonic Mean (Eq. 19)
<i>Geometric Quality</i> (Q_{geo})			Geometric Mean (Eq. 10)
$\hat{Q}_{\text{yield}}$	Measured / Authored	Ref. saturation ^a	Rank-free (Jointly constrained)
$\hat{Q}_{\text{eff}}$	Measured / Adapted	Ref. saturation ^a	
$\hat{Q}_{\text{sun}}$	Measured / Authored	Unit interval ^b	
<i>Qualitative Quality</i> (Q_{qual})			Arithmetic Mean (Eq. 11)
$\hat{Q}_{\text{green}}$	Authored / Rubric-based	Ordinal rescale ^d	Equal Weight (33.3%)
$\hat{Q}_{\text{innov}}$	Authored / Rubric-based	Ordinal rescale ^d	
$\hat{Q}_{\text{mat}}$	Authored / Rubric-based	Ordinal rescale ^d	

^a **Ref. saturation:** compresses a nonnegative raw proxy against a reference intensity r_j , producing a bounded score in $[0, 1]$ without asserting a universal empirical maximum.

^b **Unit interval:** direct continuous translation mapping spatial or orientation ratios monotonically onto $[0, 1]$.

^c **Ordinal rescale:** conversion of discrete qualitative criteria into uniform, equidistant scores bounded in $[0, 1]$.

3.7 How Height Enters the Model

Until now, the C-Q proxy framework has only developed scores in 2.5 dimensions. Yes, the X, Y, and Z axes of the parti are considered by them, but the actual costs and benefits of verticality are not — yet. While the micro and macroeconomic costs of building tall are well documented, the literature frequently overlooks how a singular geometric condition yields radically divergent feasibility outcomes depending on its vertical scale. A ten-story prismatic block and a seventeen-story tapered tower may share a clear formal lineage, yet they answer to entirely different spatial efficiency thresholds, structural demands, and submarket forces.

The construction and real estate economics literatures do not support treating height as a simple proportional surcharge. Eriksen and Orlando [2021] estimate the marginal impact of height on residential construction and show that returns to scale are nonlinear. Their analysis is useful here because it connects height to

marginal construction cost, land substitution, and the break-even rent needed to justify additional vertical construction. The exact thresholds in their study are not used as calibration points for this Seattle case, but their central finding supports the claim that height should be modeled as a distinct feasibility mechanism.

The same point appears at the building-element scale. de Jong and van Oss report an aggregate height-cost increase of roughly eight percent per ten floors, but that average hides major differences between systems. Structure, building services, and vertical transport rise steeply, elevators rise roughly 15–20% per ten floors, facade costs rise only slightly, and roof cost per unit of floor area can fall as the building grows [de Jong and van Oss, 2007]. Blackman and Picken find a similar pattern in residential towers: substructure, roof, upper floors, and door and window costs follow a U-shaped profile; services and electrical costs follow a convex profile; and plumbing rises almost linearly [Blackman and Picken, 2010]. The point at which height stops reducing unit cost and begins increasing it is also context-dependent [Flanagan and Norman, 1978, Blackman and Picken, 2010]. A single height multiplier would therefore be too blunt, because it would collapse systems that scale differently and import assumptions from another construction context.

Height is therefore embedded into the proxy system after raw geometry is extracted and normalized, but before the final transfer indices enter the feasibility model. Rather than applying one external height charge, the method reads height through the proportion of the parti. This keeps separate two questions that are often blurred: how tall the building is, and how that height is distributed over its average floor plate.

The geometric basis for this proportional reading comes from D’Amico and Pomponi [2019]. Their compactness study shows that raw (S/V) is scale-dependent and therefore cannot compare form cleanly by itself. For cuboid-like buildings, they address this by relating compactness to footprint aspect ratio and slenderness aspect ratio. This thesis does not use their compactness formula directly, because the authored schemes include podiums, setbacks, terraces, and other non-cuboid transformations. It does adopt their core lesson: height should be read together with plan proportion before it is given economic meaning.

The procedure runs in four steps. First, raw geometric proxies are extracted from the massing and normalized into comparable unitless scores. Second, a baseline-relative aspect ratio is computed to capture height and slenderness exposure. Third, the normalized scores are translated into height-exposed transfer indices. Fourth, those transfer indices enter the feasibility model through separate cost-facing, area-facing, and rate-facing channels.

Let X_0 be the baseline massing and X_i the evaluated state, with building height H_i and average floor-plate area $\bar{A}_{\text{floor},i}$. The baseline-relative aspect ratio is:

$$\chi_i = \frac{H_i / \sqrt{\bar{A}_{\text{floor},i}}}{H_0 / \sqrt{\bar{A}_{\text{floor},0}}} \quad (12)$$

At $X_i = X_0$, the ratio is $\chi_i = 1$. Values above one indicate that the scheme is taller, more slender, or both, relative to the baseline. The term $\sqrt{\bar{A}_{\text{floor},i}}$ converts average

floor-plate area into a characteristic plan dimension, so height is measured against the size of the plate carrying it. This makes χ_i a proportional exposure measure rather than a story-count variable.

This proportional reading is also consistent with the cost literature. Chau and colleagues estimate total construction cost as a joint function of height, footprint, and building quality, showing that vertical and horizontal dimensions jointly shape the economic limit of a building [Chau et al., 2007]. de Jong and van Oss make the same point from the revenue side: building efficiency, measured as gross-to-leasable area, declines with height and is weakest for tall buildings on small floor plates [de Jong and van Oss, 2007]. A slender tower is therefore charged twice, once through construction cost and once through reduced earning area. The same proportion matters structurally, since slenderness rather than height alone drives lateral demand and the systems required to resist it [Ali and Moon, 2007].

To keep measured formal complexity distinct from its height-conditioned consequence, the pre-height score C_G of Eq. (9) is exposed through the aspect ratio to form the cost-facing transfer score:

$$C_T(X_i) = \chi_i C_G(X_i). \quad (13)$$

C_T leaves the measured geometric signal intact and makes only its feasibility consequence conditional on χ_i . A taller or slenderer scheme does not rewrite its own raw measurements; it changes how forcefully those measurements bear on the model. Because the height response is carried entirely by χ_i , the downstream baseline cost is taken at the reference regime and carries no separate height term. The implied cost-versus-height behavior over the evaluated range is checked twice in the appendix, against the tall-residential rent and price elasticities of the urban-economics literature [Ahlfeldt and McMillen, 2018] and against the construction-cost slope reported for high-rise work, on the order of eight percent per ten floors, so that the internalized response neither understates nor inflates what the discourse community has measured directly [de Jong and van Oss, 2007].

The final cost-facing Complexity Index follows directly:

$$\mathbb{C}(X_i) = C_T(X_i) = \chi_i C_G(X_i). \quad (14)$$

The hierarchy is explicit. C_j is a raw complexity proxy, $\hat{C}_j$ its normalized form, C_G the aggregated pre-height geometric score, C_T the height-exposed transfer score, and $\mathbb{C}$ the final Complexity score the feasibility model reads.

Height meets the quality side asymmetrically, and the same construction-economics findings explain why. On the area side, the elements that climb fastest with height are the core-bound ones, including vertical transport and building services, and their growth often reduces the ratio of leasable to gross area [de Jong and van Oss, 2007]. The model therefore does not apply a generic height deduction. Instead, it measures the leasable area delivered by each scheme directly from the unit-area or leasable-area geometry. The measured leasable area drives the raw yield and efficiency proxies:

$$Q_{\text{yield}}(X_i) = \frac{A_{\text{lease,total}}(X_i)/F_i}{A_{\text{site}}}, \quad Q_{\text{eff}}(X_i) = \frac{A_{\text{lease,total}}(X_i)}{\sum_{f=1}^{F_i} A_{\text{floor},f}}. \quad (15)$$

These raw values are normalized into $\hat{Q}_{\text{yield}}$ and $\hat{Q}_{\text{eff}}$ via Eq. (8), then aggregated non-compensatorily into the area-facing quality transfer score:

$$Q_A(X_i) = \sqrt{\hat{Q}_{\text{yield}}(X_i) \times \hat{Q}_{\text{eff}}(X_i)}. \quad (16)$$

This places the area penalty where the evidence places it, in the amount of leasable area a massing can actually deliver. A taller or more slender scheme is not penalized simply because it is tall. It is penalized only when its height, core requirements, circulation, setbacks, or plate geometry reduce measured leasable area relative to total constructed floor area. This keeps the model consistent with de Jong and van Oss's observation that efficiency is weakest for tall buildings on small floor plates, while allowing the penalty to emerge from the authored geometry rather than from a fixed height rule [de Jong and van Oss, 2007]. When direct leasable geometry is unavailable, a core-subtraction estimate may be used as a fallback diagnostic, but it is not the primary method used in this experiment.

Rate-facing quality runs the other way. Greater height can improve daylight, outlook, and view, an effect read through the measured solar proxy and capped at unity:

$$\tilde{Q}_{\text{sun}}(X_i) = \min\left\{1, \hat{Q}_{\text{sun}}(X_i) \chi_i\right\}. \quad (17)$$

The rate-facing quality transfer score is then:

$$Q_R(X_i) = \frac{\tilde{Q}_{\text{sun}}(X_i) + Q_{\text{qual}}(X_i)}{2}, \quad (18)$$

with Q_{qual} the qualitative component from Eq. (11). The final Quality Index is the harmonic mean of the two channels:

$$\mathbb{Q}(X_i) = \frac{2 Q_A(X_i) Q_R(X_i)}{Q_A(X_i) + Q_R(X_i)}. \quad (19)$$

The quality hierarchy runs parallel to the complexity one. Q_k is a raw quality proxy, $\hat{Q}_k$ its normalized form, Q_A area-facing quality, Q_R rate-facing quality, and $\mathbb{Q}$ the final Quality Index carried into reporting and feasibility comparison.

The sequence keeps apart the effects that a single height charge would collapse, and that the literature warns against collapsing. C_G reads formal complexity and $\mathbb{C}$ its height-exposed cost consequence; Q_A reads the area penalty of carrying a growing core up a slender parti, and Q_R the rate premium of the outlook that height affords. Height is not imposed downstream as an independent escalator. It is internalized upstream through the geometry of the parti, channeled effect by effect, and calibrated to this site rather than borrowed from another, and it speaks only through

the final indices $\mathbb{C}$ and $\mathbb{Q}$, which from here on out will simply be referred to as C and Q , or (C, Q) .

From here, C and Q are integrated into the pro forma as multipliers on the cost and revenue sides respectively: C scales the hard-cost basis through the complexity premium, while Q_A and Q_R scale achievable rent or sale price through the bounded quality premiums documented in §4.5. The stochastic machinery that stresses those multipliers is taken up next.

3.8 Monte Carlo Simulation

Monte Carlo simulation (MCS) functions as a decision-support tool, not a predictive or optimizing mechanism. Its value lies not in generating a single answer, but in mapping how outcomes distribute across plausible futures and tracing how that variance shifts in response to formal design modifications [Hoesli et al., 2006].

A conventional pro forma asks a deterministic question: does this building pencil at assumed inputs? The answer is a single number. MCS shifts the inquiry to evaluate performance across a spectrum of market conditions, measuring how frequently a building clears its financial hurdles and how architectural geometry alters the profile of that risk. Consequently, two schemes with identical expected returns can harbor starkly divergent tail risks: one failing rarely but catastrophically under adverse conditions, the other maintaining a narrow, predictable variance. A point estimate conflates them; a probability distribution isolates them.

In Mode A, MCS functions as a design-space scanner, mapping probabilistic feasibility across abstract (C, Q) coordinates before any form is resolved. In Mode B, it evaluates discrete designed configurations, enabling iterative refinement through feedback on expected return, variance, and downside risk. Loizou and French [2012] argue for MCS as an exploratory device that supports sensitivity analysis and iterative reformulation of assumptions rather than a generator of definitive answers — guidance this framework operationalizes directly. French and Gabrielli [2006] and Gimpelevich [2011] extend that argument with development-feasibility applications, demonstrating that simulation reveals project sensitivities to specific input variables in ways static models do not.

Input Distributions

Each iteration draws six stochastic variables from independently specified distributions reflecting the empirical behavior of construction and real-estate data. Construction cost shock follows a bias-corrected lognormal: overruns are bounded below at zero but unbounded above, and symmetric distributions would imply equally likely large underruns, which the empirical record does not support [Touran and Wiser, 1992]; the mean is bias-corrected so the expected shock multiplier equals 1.0. Rent growth follows a clipped normal,¹ consistent with the mean-reverting

¹A clipped normal samples from $\mathcal{N}(\mu, \sigma^2)$ with rejection at fixed bounds, preserving the shape of the central mass while preventing economically implausible tail draws (e.g., negative cap rates).

treatment of income growth in stochastic DCF models [Hoesli et al., 2006]. Absorption timing follows a triangular distribution, the standard choice for variables where expert judgment bounds a plausible range but historical data is sparse [Loizou and French, 2012]. Exit capitalization rate and construction-period interest rate follow clipped normals, reflecting documented symmetric variation around prevailing yield and rate expectations [Clayton et al., 2009, Gimpelevich, 2011]. Major delay events follow a Bernoulli: most projects experience no catastrophic schedule disruption, but a meaningful minority lose substantial time in lumpy, unpredictable increments [Mayer and Somerville, 2000].

Correlation The current implementation treats the stochastic inputs as conditionally independent because this study lacks defensible market-specific covariance estimates for rent, cost, cap-rate, and financing co-movement. Correlation could be introduced in future work, but for this thesis the simplified assumption follows the portability logic of early-stage parametric risk modeling [Warszawski and Sacks, 2003]. What that means is that each variable is allowed to vary according to its own distribution, but the model does not assume that unfavorable events necessarily arrive together. This keeps the simulation transparent and portable, while acknowledging that a fully calibrated development model would eventually need to account for correlated downside conditions.

3.9 From Distributions to Decisions

A Monte Carlo iteration produces one realization of project performance under one sampled market state $\Omega^{(i)}$. That realization is *robust* when two conditions hold simultaneously:

$$\text{Robust}_i \iff \begin{cases} \text{IRR}_i \geq \tau_{\text{IRR}} & (\text{investor hurdle}) \\ \text{DSCR}_i \geq \tau_{\text{DSCR}} & (\text{lender hurdle}) \end{cases} \quad (20)$$

Both are necessary because they surface different failure modes. IRR captures profitability over the hold period but can appear strong when returns depend on back-ended exit proceeds; DSCR captures whether operating cash flow covers debt service period by period, so a project with attractive expected returns can still be non-bankable when cost escalation and rent underperformance co-occur. The non-compensatory pairing is reinforced empirically by Li and Liang [2020], whose finding that downside operating performance is disproportionately value-destructive implies that a configuration cannot be allowed to offset a debt-service failure with headline return. The reliance on a full operating-cash-flow path rather than a single capitalized figure follows the same source: Li and Liang show cap-rate-only valuation is materially too crude for heterogeneous multifamily assets under stress, and that jointly modeling market trend and relative operating performance improves predictive accuracy — direct justification for the stochastic DCF engine over a cap-rate pro forma.

Across N iterations, robustness at a design-state coordinate is the proportion satisfying both conditions:

$$P(\text{Robust} \mid C, Q) = P(\text{IRR} \geq \tau_{\text{IRR}} \cap \text{DSCR} \geq \tau_{\text{DSCR}}). \quad (21)$$

In Mode B, $(C, Q) = (C(X_i), Q(X_i))$ for a specific massing and the probability is read from the SDE surface at that coordinate. The Safe Design Envelope is the region where this probability exceeds a threshold τ_P :²

$$\text{SDE} = \{(C, Q) : P(\text{Robust} \mid C, Q) \geq \tau_P\}. \quad (22)$$

Inside the envelope a design clears its financial hurdles across at least τ_P of plausible market futures; outside it, feasibility depends on narrow assumptions, a cost escalation that stays low, an absorption schedule that holds, an exit cap rate that does not widen. The boundary between them is the object this thesis is built to find. Because it is a level set in two dimensions, it can be read along either axis: at fixed Q it returns the maximum complexity a design can carry before robustness falls below τ_P , and at fixed C it returns the minimum quality required to stay feasible. The sub-metric decomposition of $C(X_i)$ and $Q(X_i)$ then identifies which proxy pushed a non-viable scheme across the boundary, and what geometric revision would pull it back.

3.10 The Analytical Outputs

The framework treats design decisions (C, Q) as endogenous and market conditions Ω as exogenous. Project outcomes are therefore probabilistic responses to the interaction between architectural intent and environmental stress:

$$Y = f(C, Q, \Omega). \quad (23)$$

Design alternatives are evaluated not in isolation but as members of a broader set, following the tradespace exploration paradigm in systems engineering [Ross and Hastings, 2005].

Sensitivity Analysis

Before interpreting the stochastic results, a one-at-a-time sensitivity analysis identifies which inputs exert the greatest influence on project IRR. The sweep covers both the stochastic market-state vector Ω and the key fixed underwriting parameters it acts on (base rent or sale price, hard cost), since a fixed input can dominate outcome variance even when it is not itself sampled. Each input is varied across

²Mathematically, the SDE is a joint chance-constrained feasibility region in the sense of Charnes and Cooper [1959] and Prékopa [1994], transposed from financial-variable space to design-state space. The contribution is not the mathematics but the domain: no existing instrument maps a joint robustness probability across a design-state space defined by geometric complexity and asset quality. Monte Carlo sampling is the standard approximation method for such constraints [Küçükyavuz and Jiang, 2022].

its plausible range while the remaining inputs are held at their median values. The resulting change in IRR is ranked and displayed as a tornado chart, a standard feasibility tool for identifying dominant financial drivers in development analysis [Pyhrr, 1973]. In this thesis, the tornado chart is used as an interpretive check: it clarifies whether project risk is primarily driven by construction cost, rent growth, absorption, cap rates, financing, or delay assumptions.

PDFs and CDFs

For any evaluated configuration, the simulation produces a full distribution of outcomes across N iterations. The probability density function (PDF) is presented through histograms or smoothed density plots for key outputs such as IRR, DSCR, total development cost, NOI, and yield-on-cost. These plots show where outcomes concentrate, how widely they disperse, and whether the distribution is symmetric, skewed, or fat-tailed. The IRR PDF marks the hurdle rate τ alongside median and mean lines, while the DSCR PDF marks the debt-service threshold to identify the share of iterations in which the project cannot service its debt.

The cumulative distribution function (CDF) complements the PDF by showing the probability that an outcome falls at or below a given value. For IRR, the CDF reports P_{10} , P_{25} , P_{50} , P_{75} , and P_{90} percentiles, characterizing downside risk and upside potential independent of a single hurdle-rate cutoff. Together, the PDF and CDF translate probabilistic robustness into design-facing diagnostics. A PDF shifted left or spread widely suggests exposure to cost, absorption, or financing shocks. A steep CDF near the hurdle rate indicates that small changes in assumptions may determine feasibility. A DSCR breach rate elevated above the IRR failure rate directs the designer toward higher net-to-gross efficiency, lower capitalized cost, or reduced complexity premiums.

Safe Design Envelope and Efficacy ($\mathcal{E}$)

The primary Mode A output is the Safe Design Envelope, a feasibility map of the (C, Q) design space in which every coordinate has been tested against the hurdle rate across 50,000 stochastic scenarios. What the surface returns is not a single boundary line but a terrain: it shows which design combinations clear the bar, and, just as usefully, how steeply feasibility falls away as complexity ebbs and flows.

An intuitive reading would simply take the ratio Q/C , but a parti's geometric complexity and a delivered building's asset quality do not trade one-for-one. Quality lifts revenue toward a ceiling that flattens as the market's willingness to pay diminishes, while complexity increases risk on a curve that accelerates and never caps.³ To hold that asymmetry in a single figure, we introduce Efficacy, $\mathcal{E}$:

$$\mathcal{E} = \frac{\Lambda_R Q^{1/\varphi}}{\Lambda_K (1 + C^\varphi)}, \quad \varphi = \frac{1 + \sqrt{5}}{2} \approx 1.618, \quad (24)$$

³Rosen [1974] establishes the diminishing marginal willingness-to-pay that bounds the revenue channel; the construction-economics literature on the cost of formal complexity [Belniak et al., 2013, Eriksen and Orlando, 2021] documents the accelerating cost behavior of the complexity channel.

where its reciprocal, $1/\varphi \approx 0.618$, enforces the strict sublinear decay of the quality premium. Rather than carry every line of the pro forma, the economic context is compressed into two market constants, calibrated once from the baseline massing:

$$\Lambda_R = \frac{\text{NOI}_0}{\text{GFA}_0} \quad \text{and} \quad \Lambda_K = \left(\frac{\text{TDC}_0}{\text{GFA}_0} \right) \cdot \tau. \quad (25)$$

Λ_R (\$/sf/yr) is the baseline rental yield, and Λ_K (\$/sf/yr) is the baseline capitalized cost (a square foot of structure annualized through the hurdle rate τ). Both are pinned permanently to Iteration 1. Defining both constants per square foot of gross floor area makes $\mathcal{E}$ a scale-free, intensive measure. The threshold $\mathcal{E} = 1$ reads cleanly as the point where the quality-lifted yield just balances the complexity-loaded cost.

For comparison across authored iterations, a useful metric is the baseline-relative efficacy score ($\mathcal{E}_{\text{rel}}$). Because the market constants Λ_R and Λ_K are fixed, they cleanly cancel out of the ratio, leaving a direct geometric comparison:

$$\mathcal{E}_{\text{rel},i} = \frac{\mathcal{E}_i}{\mathcal{E}_0} = \left(\frac{Q_i}{Q_0} \right)^{1/\varphi} \cdot \frac{1 + C_0^\varphi}{1 + C_i^\varphi}. \quad (26)$$

This makes the interpretation direct. At the baseline ($i = 0$), $\mathcal{E}_{\text{rel},0} = 1.0$. Values above 1.0 indicate that a scheme improves the quality-to-complexity tradeoff relative to the baseline, while values below 1.0 indicate that the added complexity burden outpaces the quality gain. This relative form is essential for the Mode B sequence because the design states vary in footprint, story count, and unit count. By stripping out absolute scale, the comparison focuses on whether the design move actually improves the relationship between quality and complexity, rather than crediting a building simply for being larger, the blunt aggregate reading Li and Liang [2020] caution against.

The exponent pair $(\varphi, 1/\varphi)$ functions as a fixed theoretical control rather than a data-fitted parameter. Attempting to calculate these exact exponents using a narrow window of volatile local construction data would only create a false sense of precision. When empirical data is too noisy for reliable estimation, theory-driven weights offer a recognized and mathematically stable alternative [OECD and European Commission, Joint Research Centre, 2008]. Keeping this reciprocal pair constant allows the index to clearly illustrate the fundamental imbalance between creating value and incurring cost. The numerator $Q^{1/\varphi}$ represents diminishing returns, showing that initial investments in spatial quality yield a strong financial premium that eventually flattens as the market's willingness to pay caps out ($Q \rightarrow 1$). Conversely, the denominator $1 + C^\varphi$ models complexity as a compounding penalty. As the design moves closer to highly bespoke, non-standard geometry ($C \rightarrow 1$), the risk and cost of construction do not simply increase linearly but accelerate steeply.

3.11 Transition-Based Diagnostic Analysis

Mode B evaluates each authored scheme as a design state X_i and each design move as a transition $\Delta X_i = X_{i+1} - X_i$. This shifts the analysis from asking only whether

a scheme performs well to asking what each architectural move *changes*. Each transition records changes in complexity, quality, leasable area, cost, income, debt coverage, return, and robustness.

All schemes are simulated using common random numbers, the standard variance-reduction technique for comparing stochastic systems by holding the sampled environment fixed across configurations [Glasserman and Yao, 1992]. The same $N = 50,000$ market realizations are applied to every iteration, so paired differences isolate the effect of the design move rather than sampling noise between runs. For each realization ω_j ,

$$\Delta \text{IRR}_{i,j} = \text{IRR}(X_{i+1}, \omega_j) - \text{IRR}(X_i, \omega_j). \quad (27)$$

The same logic applies to DSCR, TDC, NOI, YOC, and robustness, so each transition can be summarized by its median effect, downside range, and probability of improvement, for example $P(\Delta \text{IRR} > 0)$. The paired construction is what permits attribution: in experimental terms it is a matched-pairs design with the sampled future as the block, so when two schemes face identical futures the difference between them is the design move, not the weather. The variance reduction is provable rather than conventional where the response is monotone in the design state, a condition Glasserman and Yao [1992] establish for exactly this class of comparison.

Applying the joint robustness condition of Eq. (20) draw by draw, each realization is then classified as robust, DSCR-only failure, IRR-only failure, or joint failure. This identifies how a scheme fails, not just whether it fails. The analysis also computes a break-even rent premium v^* , the minimum rent increase required to satisfy the governing feasibility threshold:

$$v^* = \max(v_{\text{IRR}}^*, v_{\text{DSCR}}^*). \quad (28)$$

Lower v^* values indicate closer proximity to viability; higher values indicate a scheme leaning on a larger, and potentially less realistic, market premium. Because the experiment contains five states and four transitions, the analysis stays within-project rather than claiming universal inference about types of architectural move. Drawing inferences across types would require a controlled design experiment with many variants per move type, a natural extension of this work, identified as such, but outside the present scope.

TL;DR: Designer creates parti massing in Rhino and Grasshopper then extracts C and Q . They then assemble the pro forma and run the Monte Carlo simulation in Python and Excel respectively. The Efficacy equation offers a quick deterministic screen; its elasticities try to explain how that screen behaves under pressure. The transition analysis attempts to show what actually changed across iterations. Whether it succeeds is, at this point, still an open question. The rest is architecture.

4 Project Description and Model Inputs

This chapter defines the project scenario used by the financial model. It does not evaluate the design iterations. That work starts in the following chapter. Here, the goal is simpler: explain what the Excel inputs sheet contains, why those inputs are reasonable, and how the project scenario has been bounded so the Safe Design Envelope can be tested without pretending to be a full entitlement package.

The project is a market-rate multifamily tower at 2810 Eastlake Avenue East in Seattle. It is modeled as a maritime-industrial Streamline Moderne residential tower: rounded corners, long balcony bands, smooth horizontal expression, metal railings, planted terraces, and a clear relationship to Lake Union's working-waterfront memory. The proposal is not a literal reproduction of a 1930s apartment building. It is a contemporary tower using Streamline Moderne as a formal grammar. In this case, style is not only taste. Rounded corners, balconies, terrace cuts, floor-plate efficiency, and height all become measurable design choices that later enter the simulation.

The project uses its own experimental zoning envelope. This is deliberate. The thesis is testing financial behavior under uncertainty, not the local politics of one parcel. Seattle's comprehensive planning framework is still relevant because it defines the city's broader growth logic, and Seattle's Mandatory Housing Affordability program is relevant because it shows how additional development capacity can be tied to affordable housing obligations [Seattle Office of Planning and Community Development, 2025, Seattle Department of Construction and Inspections, 2026, Seattle Office of Housing, 2026]. However, the chapter does not claim that the exact tower described here is presently entitled. The zoning envelope is a controlled research boundary. The politics are abstracted so the pro forma can stop arguing at the podium and start doing math.

4.1 Input Workbook Structure

The Excel workbook is organized as a compact input deck for the Python simulation. The workbook is not used as a black box. It provides named assumptions, reference values, and distribution parameters. Python then reads those values and recomputes the financial outputs during each simulation run.

The workbook contains four main sheets:

Table 4: Workbook sheets used to define the project scenario.

Sheet	Purpose
INPUTS	Stores the fixed project inputs: site size, gross floor area, story count, tenure switch, construction cost, revenue assumptions, financing assumptions, and robustness thresholds.
OMEGA	Stores stochastic market and delivery drivers used in the Monte Carlo simulation, including cost shock, rent or price growth, absorption, cap rate movement, interest-rate movement, and delay risk.
UNIT_MIX	Stores the unit-type structure used to derive weighted unit sizes, rents, and sale prices.
MODE_B	Stores design-iteration inputs used later in the experiment chapter. This chapter does not discuss those iteration values because the reader has not yet reached the experiment.

This structure is intentionally simple. The point is not to make the workbook look like a proprietary underwriting model. The point is to make every assumption visible enough that it can be challenged, revised, or replaced. A good research model should not require faith. It should leave its fingerprints on the glass.

4.2 Site and Planning Context

The site input begins with a real parcel and a simplified development frame. The workbook uses a site area of 16,077 square feet for 2810 Eastlake Avenue East. King County’s public parcel tools provide the official route for checking parcel, assessor, and development-condition information by address or parcel number [King County GIS Center, 2026]. Using a public parcel source matters because the site area is not a decorative assumption. It controls effective FAR, land basis per square foot, and the amount of building area that can be justified on a compact urban lot.

Eastlake is used as a experiment location because it is neither a blank slate nor a fully stabilized high-rise district. The neighborhood sits between Lake Union, Interstate 5, South Lake Union, and Capitol Hill. That condition makes the parcel useful for a height-sensitive financial experiment. Lower floors face freeway and corridor constraints. Upper floors can capture stronger light, air, and view value. The site therefore lets the model test a real development problem: whether vertical value can offset vertical cost.

The workbook does not treat contemporary Seattle zoning as a literal Day-1 entitlement analysis. Under the Seattle Municipal Code, the parcel is currently designated NC3P-55 (M) : Neighborhood Commercial 3 Pedestrian with a 55-foot height limit. That envelope would ordinarily restrict the site to a predictable low-rise podium block. Because this framework tests financial elasticities rather than permitting politics, the simulation instead assumes a High Rise (HR) zoning condition. This regulatory abstraction clears the runway for a 15-story tower with a




Figure 4: Site relative to the greater Seattle and Eastlake urban context.

baseline gross floor area of 161,822 square feet. Zoning is therefore treated as an experimental sandbox, not a permitting conclusion. The model is not asking whether this project would pass design review next month; it is testing how a tower-like parti behaves when cost, revenue, financing, and design quality are exposed to market uncertainty.



Figure 5: Close-in site plan. North is up.

Table 5: Primary site and program inputs from the `INPUTS` sheet.

Input	Workbook value	Use in the model
Site area	16,077 sf	Parcel denominator for effective FAR and land basis
Gross floor area	161,822 sf	Baseline project scale
Story count	15	Baseline tower height used before design iterations
Dwelling units	150	Based on ten units per occupied level
Average unit size input	807 sf	Based on a unit mix of five 1-BR, three 2-BR, and two 3-BR units per level.
Tenure switch	1 = rental, 2 = condo	Allows the same physical project to be tested as either apartment or for-sale housing

4.3 Architectural Scenario

The architectural scenario is a maritime-industrial Streamline Moderne multifamily tower. Streamline Moderne is useful here because its formal logic is legible and measurable. DOCOMOMO US describes the style as strongly horizontal, defined by curved surfaces, curved corners, cylindrical projections, and restrained ornament [DOCOMOMO US, 2026]. SAH Archipedia similarly identifies horizontal lines, rounded corners, and occasional nautical elements such as railings and porthole-window motifs [SAH Archipedia, 2026]. Some of these traits, like schematic massing elements are directly translatable into the design variables described in the previous chapter, while others like windows and railings live further along the development process beyond the scope of the experiment.

Tamarkin Co.'s 10 Sullivan tower in Manhattan is a useful contemporary precedent: an expressive, non-prismatic geometry that converts a difficult urban condition into a premium spatial asset [Chaban, 2016, Rong et al., 2020]. Sited on a tapering wedge beside heavy traffic infrastructure, the sixteen-story project uses a rounded southern prow and corbelled brick to join industrial memory with a streamlined Modernism [Tamarkin Company, 2026]. Rather than retreating to a generic footprint, the design treats the constraint as the source of its identity, and the curved glazing wrapping the prow widens interior panoramic views. The precedent matters here for a narrow reason: it shows formal articulation producing institutional real estate value on a constrained site, which is the proposition the experiment is built to measure.

Because 10 Sullivan is positioned unapologetically at the upper echelon of the New York City market-rate condominium sector, this Seattle experiment may chart a similar economic trajectory as feasibility is systematically explored. This baseline positioning is not the result of its architectural form per se, but is rather driven by the raw mass and scale of the underlying building envelope, which establishes the absolute financial capacity of the site before localized geometric modifications are

ever introduced. The macroscopic dimensions of the tower define the volumetric capacity available to distribute a fixed land basis and heavy Type I-A structural overhead across the project.



Figure 6: Images found at Tamarkin Co., “10 Sullivan,” accessed June 10, 2026, <http://tamarkinco.com/10-sullivan>.

The important point is that the architectural language is chosen before the experiment begins. The design is not reverse-engineered from the best-performing coordinate. By treating the formal grammar as an independent variable rather than an optimized computational output, the simulation can cleanly isolate the distinct economic elasticities of specific geometric moves. The model will later test whether the chosen language can carry its own financial weight.

4.4 Construction Cost Inputs

The baseline construction inputs are shown in Table 6. The workbook uses a baseline hard cost of \$420 per gross square foot, soft costs equal to 22% of hard costs, a 5% contingency, and a \$6.0 million land cost. The hard cost of \$420 may seem low for a scheme like this, but as discussed earlier, the C ad Q system impact this figure with scale.

Table 6: Construction and land-cost inputs from the `INPUTS` sheet.

Input	Workbook value	Justification logic
Hard cost	\$420/sf	Local tower-scale multifamily benchmark, cross-checked against public Seattle construction-cost trend sources
Soft cost rate	22%	Development-cost allowance for design, engineering, permits, fees, financing, legal, and other non-hard-cost items
Contingency	5%	Base reserve, pre stochastic cost shock
Land cost	\$6,000,000	Fixed acquisition assumption based on the 2026 King County Appraised value of \$5.6m.

Additionally, the hard-cost value is not treated as certain. It is only the center of the cost model. Public cost sources are used to support the order of magnitude and the local escalation context. Mortenson’s Seattle Construction Cost Index is useful because it tracks representative construction pricing in the Seattle market on a quarterly basis [Mortenson, 2026]. To reinforce this framework, the Rider Levett Bucknall (RLB) North American Quarterly Construction Cost Report provides a vital independent institutional benchmark, tracking specialized high-rise delivery metrics, regional labor capacity constraints, and localized structural commodity escalation [Rider Levett Bucknall, 2026]. Utilizing dual institutional data streams confirms that the model’s central hard-cost assumptions remain resilient against regional construction cost inflation anomalies. The 15-story scale also matters for construction type. The International Building Code (IBC) links allowable height and area to occupancy and construction type, and Chapter 6 defines the construction-type framework used for fire-resistance logic [International Code Council, 2024, 2021].

The ICC also explains the high-rise threshold as an occupied floor more than 75 feet above the lowest level of fire department vehicle access [International Code Council, 2019]. Breaching this regulatory boundary forces a fundamental typological shift from standard wood framing over a concrete podium to an exclusively Type I noncombustible structural assembly. This transition mandates a complex array of active and passive life-safety infrastructure, including pressurized exit enclosures, mechanical smoke control systems, heightened hourly structural fireproofing, and even insurance considerations. Obviously, a complete code analysis is omitted at the abstract level of this experiment, but it is worth mentioning the threshold as a necessary physical justification for proxied non-linear cost step-functions applied within the C-Q aggregation to capture vertical engineering penalties to a degree.

4.5 Revenue and Tenure Inputs

The workbook can run the project in two tenure modes. Apartment mode treats the project as a standard rental multifamily building. Condo mode treats it as a for-sale condominium project. The base setting is Apartment mode but both sets of inputs

are retained because the same architectural mass can produce different risk profiles under different tenure assumptions.

Table 7: Revenue inputs from the `INPUTS` and `UNIT_MIX` sheets.

Input	Workbook value	Use in the model
Apartment rent	\$5.00/sf/mo	Baseline rental revenue before quality premium and stochastic growth
Quality rent premium	15% at full quality	Upper bound for design-driven rent uplift in rental mode
Vacancy rate	5%	Stabilized vacancy deduction in rental mode
Operating expense ratio	22%	Operating expenses as a share of effective gross income
Apartment absorption	18 months	Lease-up or build-duration input for rental timing
Condo sale price	\$950/sf	Baseline for-sale revenue before quality premium and stochastic price drift
Condo quality premium	20% at full quality	Upper bound for design-driven sale-price uplift in condo mode
Sellout period	24 months	Condo sellout duration after delivery
Sales cost	6%	Commissions and marketing as a share of gross sales
Condo margin floor	15%	Minimum profit margin used in condo robustness testing

The rent and sale-price assumptions are intentionally exposed rather than buried. Public market reports are imperfect, but they are useful for anchoring the model to observable market conditions. Cushman & Wakefield’s Seattle multifamily reports provide public information on rents, vacancy, employment, and pipeline conditions for the regional apartment market [Cushman & Wakefield, 2025]. Yardi Matrix provides another public cross-check on Seattle rent and occupancy conditions [Yardi Matrix, 2026]. These sources do not prove that a specific Eastlake tower will achieve the exact workbook rent. They establish that the input sits inside a plausible Seattle multifamily range, which is the correct standard for a thesis experiment.

To ground these figures specifically, the model’s baseline rental rate of \$5.00 per square foot and sales basis of \$950 per square foot reflect the upper tier of the contemporary Class A urban core supply. Cushman & Wakefield’s Q4 2025 metrics highlight stable absorption within the contiguous South Lake Union and Eastlake submarkets, anchored by proximity to localized tech and life sciences employment clusters. Concurrently, the February 2026 Yardi Matrix data demonstrates that while city-wide trailing averages remain moderate, premium urban high-rise configurations command clear pricing power as the construction pipeline for com-

plex downtown towers begins to thin. Calibrating the baseline to these verified institutional underwriting boundaries ensures the simulation avoids arbitrary optimization. Instead, it positions the initial project state at a historically defensible starting point, guaranteeing that any financial degradation tracked during the subsequent parametric variations represents an authentic risk signature of the building’s geometry rather than an artifact of speculative underwriting.

The `UNIT_MIX` sheet refines the revenue model by separating unit types. The current mix contains no studios, 50% one-bedroom units, 30% two-bedroom units, and 20% three-bedroom units. This produces a weighted average unit size of approximately 808 square feet in the reference mix.

Table 8: Reference unit mix from the `UNIT_MIX` sheet.

Unit type	Share	Units	Avg. sf	Sale price
Studio	0%	0	0	\$0/sf
One-bedroom	50%	75	552	\$980/sf
Two-bedroom	30%	45	1,002	\$920/sf
Three-bedroom	20%	30	1,156	\$880/sf
Total / weighted	100%	150	808	Model-derived

This mix matters because larger units increase the importance of private outdoor space, view quality, plan efficiency, and amenity logic. A balcony is not just a line item. In this model, it is part of the argument over whether quality can earn its keep.

4.6 Financing Inputs and Robustness Thresholds

The financing inputs establish the capital structure and the minimum performance thresholds used by the robustness test.

Table 9: Financing and threshold inputs from the INPUTS sheet.

Input	Value	Use in the Model
Loan-to-cost	65% ¹	Debt share of total project cost
Loan rate	7.0%	Baseline borrowing rate before stochastic rate shock
Amortization	30 years ²	Rental-mode permanent debt amortization
DSCR threshold	1.25x ³	Debt-service coverage floor for rental robustness
IRR hurdle	12% ⁴	Minimum return threshold
Hold period	7 years ⁵	Apartment-mode hold period before disposition
Exit cap rate	5.0% ⁶	Stabilized NOI capitalization rate in rental mode
Cost of sale	2.5%	Disposition cost deduction in rental mode

¹ See [Brueggeman and Fisher, 2021]. A 60–70% Loan-to-Cost baseline represents standard commercial and multifamily underwriting parameters to limit lender default exposure while ensuring sufficient developer equity cushion under modern banking prudential frameworks.

² See [Linneman, 2018]. A 30-year amortization schedule is standard for permanent institutional multifamily debt assets, matching typical agency loan programs (e.g., Fannie Mae and Freddie Mac) to preserve early-stage cash-flow yield.

³ Explained and justified in §3.2.

⁴ See [Moorhead et al., 2024]. Based on a survey of 225 real estate developers, they demonstrate that IRR is the dominant go/no-go metric in development feasibility analysis, with hurdle rates commonly ranging from 12–20% for market-rate multifamily development depending on risk profile and market context.

⁵ See [Geltner et al., 2014]. Under institutional real estate underwriting paradigms, a 5-to-10-year holding period (with 7 years serving as the modal industry benchmark) is standard for calculating multi-year pro forma cash flows before executing capital recycling.

⁶ See [Ling and Archer, 2020]. Institutional financial underwriting models routinely mandate an exit cap rate premium expansion (typically 25 to 50 basis points over the terminal or initial capitalization rate) to account for structural asset aging, physical depreciation, and long-term macro market fluctuations.

4.7 Stochastic Inputs

The OMEGA sheet stores the stochastic variables sampled during the Monte Carlo simulation. These variables represent uncertainty in the market and delivery environment. The thesis uses Monte Carlo simulation because real estate development is not governed by one clean forecast. It is shaped by ranges: cost overruns, price drift, absorption speed, interest-rate movement, cap-rate movement, and delay events. Loizou and French argue that Monte Carlo simulation is useful in real estate development because it forces the analyst to evaluate risk as a distribution rather than as a single deterministic outcome [Loizou and French, 2012].

Table 10: Stochastic market vector Ω . These are the only quantities sampled.

Variable	Distribution	Parameters
Construction cost shock	Lognormal, mean-cor.	Unit: multiplier Mean: 1.00 Spread: $\sigma \approx 0.12$ Bounds: [0.80, 1.40]
Rent growth	Normal, clipped	Unit: %/yr Mean: 2.5 Spread: $\sigma = 1.5$ Bounds: [-3, 8]
Absorption scalar	Triangular	Unit: multiplier Mode: 1.00 Bounds: [0.70, 1.40]
Exit cap rate	Normal, clipped	Unit: % Mean: 5.0 Spread: $\sigma = 0.8$ Bounds: [3.5, 8.5]
Construction Intr. rate	Normal, clipped	Unit: % Mean: 7.0 Spread: $\sigma = 1.0$ Bounds: [4.0, 12.0]
Major delay event	Bernoulli	Probability: $p = 0.08$ Duration: 3–6 months, uniform Outcome: $\{0, 1\}$

The distributions are intentionally modest. The model is not trying to simulate the apocalypse, nor is it pretending the market is a calm pond. The selected ranges create enough turbulence to test robustness while remaining legible. A project should not pass because the weather is always perfect. It should pass because it can take a few waves over the bow and still make harbor.

4.8 Design Variables Reserved for the Experiment

The workbook includes design variables for Complexity, C , and Quality, Q , but this chapter does not evaluate any parti or design iteration. As made crystal clear by this point, Complexity, C , is the normalized cost-side design index. It represents the degree to which a design departs from a simple, efficient baseline through additional envelope burden, formal articulation, geometric variation, structural complication, or other construction-relevant features. Quality, Q , is the normalized value-side design index. It represents the degree to which a design improves market appeal through better views, balconies, daylight, amenity, identity, and overall livability. The next chapter will illustrate authored design iterations and their extracted C and Q values.

5 The Experiment

The methodology developed in Chapter 3 is a two-mode framework, and its claim to operability depends on Mode B working as advertised. Mode A maps the abstract feasibility region; Mode B authors specific architectural configurations within (or against) that region and asks whether the framework’s diagnostics can price each successive design revision as the design develops. This chapter is the log of that experiment. Each iteration records the architectural change made, the resulting C and Q index values extracted from the Grasshopper geometry, the full diagnostic readout from the Monte Carlo engine, and a short commentary on what the diagnostic revealed and what change it motivated next. The format is deliberate. The chapter reads as a laboratory notebook rather than a narrative, because the thesis argument requires that the framework guide design action as opposed to scoring finished schemes.

One point of method governs the entire sequence and should be stated plainly at the outset. The five schemes were authored as architecture first. Each move follows from a design, programmatic, or experiential intention, not from a search for the bottom line. The schemes were not re-drawn to chase feasibility after the fact. The role of the model is to make the financial consequence of each architectural decision legible, so the designer can see what a given gesture costs before the design hardens.

Building on the Eastlake corridor calibration detailed in Chapter 4, the experiment isolates geometric performance by holding all baseline underwriting parameters constant across a 15- to 24-story vertical testing sequence. The model assumes a 16,000 SF parcel upzoned to an HR zone with undefined FAR limits, utilizes a baseline hard cost of \$420 per square foot, and a conventional joint 12% IRR and $1.25 \times$ DSCR threshold. Each architectural iteration is subjected to $N = 50,000$ stochastic simulation draws to generate a statistically healthy dataset. Rather than evaluating absolute financial yields in isolation, the analysis focuses on the divergence between the baseline-relative Financial Efficacy index ($\mathcal{E}_{\text{rel}}$) and the joint robustness probability ($P(\text{Robust} \mid C, Q)$). While efficacy measures how efficiently a given parti converts geometric complexity into asset quality within the static Mode A surface, robustness tracks whether the resulting pro forma survives capital market uncertainty. Observing these two indicators decouple under baseline rental tenure, followed by a comparative cross-check under a for-sale condominium structure, establishes the primary analytical jazz of this chapter.

Iteration Template

For comparability, every iteration follows the same six-part structure⁴:

1. **Architectural move.** The single design change introduced relative to the prior iteration, described as the parti revision the designer is committing to.

⁴The iteration sequence below is restricted to architectural moves operating within *conventional construction technology*: cast-in-place concrete structure, standard floor-to-floor heights, stick-built or unitized envelopes, and ordinary core-and-shell delivery.

2. **Geometry.** The Rhino model state, summarized through the key spatial variables (story count, plate dimensions, perimeter condition, core position, height, and any non-default features).
3. **Index extraction.** The C and Q values returned by the Grasshopper pipeline, with sub-metric decomposition where the move targeted a specific component.
4. **Diagnostic readout.** The standardized Monte Carlo output: median IRR with P10/P90 envelope, σ_{IRR} , $P(\text{IRR} \geq \tau)$, median DSCR, $P(\text{Robust})$, YOC, TDC, NOI, and the efficacy index.
5. **Reading.** A short commentary on what the diagnostic revealed: which constraint was binding, how the efficacy and robustness lenses agreed or disagreed, and where the scheme sat relative to the SDE boundary.
6. **Move forward.** The architectural revision motivated by the reading, which becomes the architectural move of the next iteration.

5.1 Iteration 1 – Baseline

Architectural move A trapezoidal extrusion of the maximum buildable footprint, taken from the setback lines and carried to fifteen stories above a single-story ground-floor podium, with a centered double-loaded core. Each plate is uniform and accommodates ten stacked units, for 150 units total. The intent is to establish the baseline $P(\text{Robust})$ value against which every later architectural move can be measured. This is the most basic massing in the set, and it is included precisely because it spends nothing on expression at this resolution.




Figure 7: The Least Suspicious Box

Geometry

Stories	15
Setbacks	10' front/back, 15' sides
Site area	16,078 sf
Gross floor area	161,822 sf
Leasable floor area	121,050 sf
Typical plate	≈ 10,788 sf/floor
Typical leasable area	≈ 8,070 sf/floor
Core	≈ 946 sf/floor
Units	150




Figure 8: Aerial view, Iteration 1.




Figure 9: Floor plan, Iteration 1.

Index extraction $C = 0.162$, $Q = 0.452$. Complexity is low and almost entirely envelope-driven: the uniform stack produces zero vertical and edge variation

($C_{\text{stack}} = C_{\text{edge}} = 0$), so the only contributors are the surface-to-volume burden of the skin and a modest core offset. With no height premium, $\chi = 1.000$ and $C = C_G$. On the quality side, the baseline is the calibration anchor for the reference set $\mathcal{R}$: its area-side scores are pinned to $\hat{Q}_{\text{yield}} = \hat{Q}_{\text{eff}} = 0.500$ by construction (the raw efficiency score of 0.91 and raw yield of 0.61 both calibrate to 0.50), so the building reads as moderate quality, held down by a low qualitative score of 0.300 on a scheme that has not yet committed to any expressive material or amenity program.

Iteration 1 – Trapezoidal Baseline

<i>Architectural Complexity</i>	Value
Envelope, C_{env}	0.040
Stack variation, C_{stack}	0.000
Edge variation, C_{edge}	0.000
Core offset, C_{core}	0.065
Height multiplier, χ	1.000
Pre-height geometric complexity, C_G	0.162
Composite Architectural Complexity, C	0.162

<i>Asset Quality</i>	Value
Area yield, $\hat{Q}_{\text{yield}}$	0.500
Area efficiency, $\hat{Q}_{\text{eff}}$	0.500
Solar exposure, Q_{sun}	0.528
Qualitative, Q_{qual}	0.300
Area-side composite, Q_A	0.500
Rate-side composite, Q_R	0.414
Composite Asset Quality, Q	0.452

Complexity rows report raw geometric readings (e.g., surface-to-volume ratio C_{env}) normalized via §3.6 reference-saturation before rank-decay aggregation, rather than combining arithmetically into C_G . Quality rows reflect post-calibration scores. Final complexity is $C = C_G \cdot \chi$; final quality is the harmonic mean of Q_A (geometric mean of yield and efficiency scores) and Q_R (mean of the qualitative and height-adjusted solar score $\tilde{Q}_{\text{sun}} = \min\{1, \hat{Q}_{\text{sun}} \chi\}$).

Diagnostic readout

Iteration 1 – Trapezoidal Baseline	GFA: 161,822 sf
Median IRR	11.7%
P10 / P90 IRR	0.7% / 20.5%
σ_{IRR}	7.7%
$P(\text{IRR} \geq \tau)$	49%
Median DSCR	1.22×
$P(\text{Robust} \mid C, Q)$	35%
Yield-on-Cost (YOC)	6.32%
Total Development Cost (TDC)	\$100.2M
NOI (stabilized)	\$6.34M
Efficacy, $\mathcal{E}$ ($\mathcal{E}_{\text{rel}}$)	0.307 (1.000)




Figure 10: Iteration 1 - Baseline

Reading The baseline is the strongest scheme in the sequence and is still not quite feasible, which is a humbling way to open a thesis about architectural ambition. Median IRR is 11.7%, sitting just under the $\tau = 12\%$ hurdle, and median DSCR is $1.22\times$, just under the $1.25\times$ floor. Just under half of the draws clear the IRR hurdle ($P(IRR \geq \tau) = 49\%$), but only 35% are jointly robust. That fourteen-point gap is the debt-service constraint doing its work: a meaningful share of the scenarios that earn an acceptable return still fail to cover debt at $1.25\times$. The scheme therefore sits just outside the Safe Design Envelope, and it does so not because of any formal excess but because the base development condition is marginal under ordinary market uncertainty.

This matters for everything that follows. The trapezoidal extrusion is the most efficient move available. It maximizes the setback footprint, holds a uniform fifteen-story stack, and carries the highest leasable area of the early schemes at 121,050 sf. Its complexity is low at $C = 0.162$ and its quality moderate at $Q = 0.452$, so the project is not being penalized for ambition it has not yet shown. The constraint is revenue capacity relative to capital. With the baseline pinned just below both thresholds, each subsequent move starts from a position where there is very little slack to spend. Balconies, a sculpted corner, a rooftop amenity, and added height are not decorations laid over a comfortable pro forma; they are gestures that have to be paid for out of a budget that is already empty.

Move forward With the baseline fixed, the sequence begins to spend. Iteration 2 introduces perimeter setback balconies, trading rentable area for exterior amenity and a higher qualitative score. The question is not whether balconies reduce efficiency, which they plainly do, but whether the quality they add is cheap enough in complexity terms to register as a defensible move once the efficacy and robustness lenses are read together.

5.2 Iteration 2 – Perimeter Setback Balconies

Architectural move Iteration 2 steps the upper floors back by five feet around the perimeter to carve exterior balconies into floors 5 through 15. The lower floors retain the Iteration 1 plate, while the upper floors shift to a reduced unit count to hold the balcony zone and preserve workable unit dimensions. A few units on the east and west faces keep the full footprint where the plates are already tight. The move is amenity led: it asks whether private outdoor space and a more articulated facade are worth the rentable area they consume.




Figure 11: Great Views, Bad News

Geometry

Stories	15
Setbacks	10' front/back, 15' sides
Balcony setback	5' perimeter setback on floors 5–15
Site area	16,078 sf
Gross floor area	161,822 sf
Leasable floor area	107,168 sf
Typical lower plate	≈ 10,788 sf/floor
Typical leasable area	≈ 7,145 sf/floor
Core	≈ 946 sf/floor
Units	120




Figure 12: Aerial view, Iteration 2.




Figure 13: Floor plan, Iteration 2.

Index extraction $C = 0.165$, $Q = 0.475$. Complexity barely moves. Floor area, perimeter, height, stack variation, and edge variation are all unchanged from the baseline, so the only shift is a slight rise in the envelope score as the balcony recesses add skin per unit volume; χ remains 1.000. Quality rises, and it rises on the rate side: the qualitative score climbs from 0.300 to 0.567 as the balconies register as a genuine amenity, lifting Q_R to 0.547. The area side, by contrast, falls. The balcony setback removes interior floor area, so Q_A drops from 0.500 to 0.421 and leasable area contracts from 121,050 to 107,168 sf, with the unit count falling from 150 to 120. The composite Q still improves, because the rate-side gain outweighs the area-side loss in the quality index, but the loss is real and the pro forma will find it.

Iteration 2 – Perimeter Setback Balconies

<i>Architectural Complexity</i>	Value
Envelope, C_{env}	0.043
Stack variation, C_{stack}	0.000
Edge variation, C_{edge}	0.000
Core offset, C_{core}	0.065
Height multiplier, χ	1.000
Pre-height geometric complexity, C_G	0.165
Composite Architectural Complexity, C	0.165
<i>Asset Quality</i>	Value
Area yield, $\hat{Q}_{\text{yield}}$	0.421
Area efficiency, $\hat{Q}_{\text{eff}}$	0.421
Solar exposure, Q_{sun}	0.528
Qualitative, Q_{qual}	0.567
Area-side composite, Q_A	0.421
Rate-side composite, Q_R	0.547
Composite Asset Quality, Q	0.475

Diagnostic readout

Iteration 2 – Perimeter Setback Balconies	GFA: 161,822 sf
Median IRR	7.0%
P10 / P90 IRR	-6.1% / 16.7%
σ_{IRR}	9.1%
$P(\text{IRR} \geq \tau)$	26%
Median DSCR	1.08×
$P(\text{Robust} \mid C, Q)$	11%
Yield-on-Cost (YOC)	5.61%
Total Development Cost (TDC)	\$100.3M
NOI (stabilized)	\$5.64M
Efficacy, $\mathcal{E}$ ($\mathcal{E}_{\text{rel}}$)	0.316 (1.030)




Figure 14: IRR distribution for Iteration 2, with the baseline shown for reference (grey). The distribution shifts left of the baseline while remaining tighter on the upside.

Reading This is the iteration where the two lenses part company, and the divergence is the most useful single result in the chapter. The pro forma weakens sharply: median IRR falls from 11.7% to 7.0%, the share of draws clearing the hurdle drops from 49% to 26%, and joint robustness collapses from 35% to 11%. Median DSCR slips to $1.08\times$. Yet the efficacy index rises, from $\mathcal{E}_{\text{rel}} = 1.000$ to 1.030. By the efficacy reading, the balconies are a good move; by the robustness reading, they are a costly one.

The reconciliation is precise. Efficacy measures how much quality a geometry buys per unit of complexity, and balconies buy quality cheaply: Q rises while C barely moves from 0.162 to 0.165. Robustness measures whether the cash flow survives, and the balconies are paid for almost entirely in lost leasable area, not in complexity or capital. Total development cost is essentially flat, from \$100.2M to \$100.3M, but leasable area falls by nearly 14,000 sf and the unit count drops by thirty. NOI falls to \$5.64M and YOC to 5.61%. The penalty is therefore a revenue penalty, and because it does not touch the complexity axis, it is invisible to the efficacy surface and fully visible to the Monte Carlo engine. The framework localizes the cost: the balconies are an efficient way to add architectural quality and an inefficient way to underwrite a rental pro forma, and those are not contradictory statements. They are the same building seen along two different axes.

Move forward Iteration 3 does not try to recover the lost area. It pushes the articulation further, replacing the repetitive perimeter balcony with a single sculpted gesture at the prominent southwest corner. The motivation is identity rather than yield: the corner becomes the public face of the project and consolidates the Streamline Moderne language the tower is developing. Because the move trades a diffuse balcony loss for a concentrated curved one, it is also a test of whether the framework can tell a focused expressive gesture apart from a distributed one.

5.3 Iteration 3 – Swept Southwest Corner

Architectural move Iteration 3 introduces a large swept curve at the southwest corner, the most visible corner of the site. Rather than treating the perimeter as a repetitive balcony operation, this scheme concentrates the gesture, smoothing the corner, strengthening the horizontal expression, and giving the tower a deliberate public-facing profile. The move is intuitive rather than pro forma driven.

Importantly, because the 89ft building footprint already operates at its relative physical limit within the site width, the envelope is bound by the imposed site-line setbacks. Appending traditional accessory balconies to the exterior facade would illegitimately breach the ground setback buffer. To maintain the tectonic elegance of a unified floor plate system, the balcony operation must carve inward rather than project outward, cannibalizing interior unit area as a consequence. While future iterations might test the regulatory legality of treating these recessed zones as a specialized class of leasable space, this framework treats them as an unmonetized loss—a values test where the model is asked to price a decision that architectural judgment already considers worthwhile.




Figure 15: Curve Appeal

Geometry

Stories	15
Setbacks	10' front/back, 15' sides
Corner condition	5' perimeter setback with swept southwest corner
Site area	16,078 sf
Gross floor area	159,204 sf
Leasable floor area	103,774 sf
Typical plate	≈ 10,614 sf/floor
Typical leasable area	≈ 6,918 sf/floor
Core	≈ 946 sf/floor
Units	120




Figure 16: Aerial view, Iteration 3.




Figure 17: Floor plan showing the swept southwest corner and balcony system, Iteration 3.

Index extraction $C = 0.155$, $Q = 0.471$. Complexity actually falls, which is the counterintuitive result the new metric is built to capture. The sweep regularizes the plan: it trims perimeter and reduces the core-offset score, and although it introduces a slight height premium ($\chi = 1.008$), the net pre-height complexity drops to 0.155. A curved corner reads to the eye as more elaborate than a square one, but to the geometry it is smoother, and the index reports the geometry. Quality holds just below Iteration 2. The rate side stays elevated on a strong qualitative score of 0.567, while the area side falls again as leasable area drops to 103,774 sf, the lowest of the fifteen-story schemes. The swept corner improves the building's character without improving its revenue geometry.

Iteration 3 – Swept Southwest Corner	
<i>Architectural Complexity</i>	Value
Envelope, C_{env}	0.042
Stack variation, C_{stack}	0.000
Edge variation, C_{edge}	0.000
Core offset, C_{core}	0.053
Height multiplier, χ	1.008
Pre-height geometric complexity, C_G	0.155
Composite Architectural Complexity, C	0.155
<i>Asset Quality</i>	Value
Area yield, $\hat{Q}_{\text{yield}}$	0.413
Area efficiency, $\hat{Q}_{\text{eff}}$	0.417
Solar exposure, Q_{sun}	0.525
Qualitative, Q_{qual}	0.567
Area-side composite, Q_A	0.415
Rate-side composite, Q_R	0.546
Composite Asset Quality, Q	0.471

Diagnostic readout

Iteration 3 – Swept Southwest Corner		GFA: 159,204 sf
Median IRR		6.3%
P10 / P90 IRR		-7.2% / 16.1%
σ_{IRR}		9.3%
$P(\text{IRR} \geq \tau)$		24%
Median DSCR		1.06×
$P(\text{Robust} \mid C, Q)$		9%
Yield-on-Cost (YOC)		5.52%
Total Development Cost (TDC)		\$98.6M
NOI (stabilized)		\$5.45M
Efficacy, $\mathcal{E}$ ($\mathcal{E}_{\text{rel}}$)		0.316 (1.029)




Figure 18: IRR distribution for Iteration 3, baseline shown for reference (grey).

Reading Iteration 3 confirms the pattern Iteration 2 introduced and adds a wrinkle. The pro forma drifts down again: median IRR to 6.3%, hurdle clearance to 24%, robustness to 9%, median DSCR to $1.06\times$. Efficacy holds essentially flat at $\mathcal{E}_{\text{rel}} = 1.029$, still above the baseline. The swept corner is, by the efficacy reading, still a defensible move, and it is also the cheapest scheme in the set: at \$98.6M it carries the lowest total development cost, because the sweep removes a little gross area along with the rentable area. The wrinkle is that this is the scheme where reduced complexity and reduced cost coincide with reduced feasibility, which only makes sense once the two lenses are kept separate. The geometry got simpler and cheaper, and the revenue base got smaller, and in a rental pro forma the second effect dominates. The result is again a measurement rather than a verdict: the corner is worth roughly what it costs in lost area, which is most of the additional distance to the hurdle, and the framework has now priced two consecutive expression-led moves whose efficacy stays positive while their robustness erodes.

Move forward The first three schemes hold height constant and lose area. Iteration 4 changes register. It adds program rather than subtracting it, inserting a rooftop amenity deck and two larger units per floor on the upper levels, and in doing so it adds height for the first time. The move tests whether shared amenity and a stronger upper-level silhouette can carry value the way private balconies did, and it is the first scheme to put the height multiplier χ into play.

5.4 Iteration 4 – Rooftop Amenity Deck

Architectural move Iteration 4 introduces a rooftop amenity deck and two larger units per floor on the upper levels, taking the tower to seventeen stories. The amenity reads as a prominent upper-level stepback, strengthening the horizontal massing and making the building more legible as a Streamline Moderne tower. The intent is programmatic generosity, not rescue. There is no direct pro forma line for a shared roof deck, so its value enters through the qualitative score Q_{qual} , which the designer raises to reflect the added amenity and ambition. The iteration tests how much of that intention the quality framework can carry while the pro forma keeps responding mainly to area, cost, and coverage.




Figure 19: Decked Out

Geometry

Stories	17
Approx. height	174 ft
Setbacks	10' front/back, 15' sides
Massing	Balcony system with rooftop amenity stepback
Site area	16,078 sf
Gross floor area	173,970 sf
Leasable floor area	109,810 sf
Typical plate	≈ 10,234 sf/floor
Typical leasable area	≈ 6,459 sf/floor
Core	≈ 946 sf/floor
Units	124




Figure 20: Aerial view, Iteration 4.




Figure 21: Floor plan, Iteration 4.

Index extraction $C = 0.387$, $Q = 0.495$. Complexity more than doubles, and the decomposition shows exactly why. The sectional move activates stack variation for the first time ($C_{\text{stack}} = 0.106$) as the upper plates change, introduces edge variation ($C_{\text{edge}} = 0.020$), and pushes the height premium to $\chi = 1.193$. The pre-height complexity of 0.325 is then amplified by height to a final $C = 0.387$. Quality climbs to its second-highest value, $Q = 0.495$, on a strong qualitative score of 0.667 and a high rate side ($Q_R = 0.644$). The area side stays weak, however: average leasable area per floor falls to roughly 6,459 sf as the amenity displaces rentable plate, so Q_A holds near 0.40. The scheme is more ambitious and more complex at once, and for the first time the complexity is growing faster than the quality.

Iteration 4 – Rooftop Amenity Deck

<i>Architectural Complexity</i>	Value
Envelope, C_{env}	0.044
Stack variation, C_{stack}	0.106
Edge variation, C_{edge}	0.020
Core offset, C_{core}	0.063
Height multiplier, χ	1.193
Pre-height geometric complexity, C_G	0.325
Composite Architectural Complexity, C	0.387
<i>Asset Quality</i>	Value
Area yield, $\hat{Q}_{\text{yield}}$	0.396
Area efficiency, $\hat{Q}_{\text{eff}}$	0.408
Solar exposure, Q_{sun}	0.520
Qualitative, Q_{qual}	0.667
Area-side composite, Q_A	0.403
Rate-side composite, Q_R	0.644
Composite Asset Quality, Q	0.495

Diagnostic readout

Iteration 4 – Rooftop Amenity Deck	GFA: 173,970 sf
Median IRR	1.6%
P10 / P90 IRR	-13.4% / 12.7%
σ_{IRR}	10.0%
$P(\text{IRR} \geq \tau)$	12%
Median DSCR	0.96×
$P(\text{Robust} \mid C, Q)$	2%
Yield-on-Cost (YOC)	5.00%
Total Development Cost (TDC)	\$113.0M
NOI (stabilized)	\$5.65M
Efficacy, $\mathcal{E}$ ($\mathcal{E}_{\text{rel}}$)	0.281 (0.916)




Figure 22: IRR distribution for Iteration 4, baseline shown for reference (grey). The median sits barely above zero and the lower tail has spread well past it.

Reading Iteration 4 is the turn. For the first time the two lenses agree, and they agree that the move is not worth it. Efficacy drops below the baseline to $\mathcal{E}_{\text{rel}} = 0.916$, because the complexity surge from the sectional move and the height premium has finally outrun the quality gain. The pro forma drops with it: median IRR to 1.6%, hurdle clearance to 12%, robustness to 2%. The binding constraint has also shifted. Median DSCR falls to $0.96\times$, below 1.0, meaning that in the central case the stabilized NOI no longer covers debt service at all, let alone at $1.25\times$. Where the baseline failed mainly on return, this scheme fails first on coverage. Capital is the reason: total development cost jumps to \$113.0M as height and section drive both the structure and the envelope, while NOI recovers only to \$5.65M on the slightly larger floor count. The amenity deck is a real improvement to the building and a clear loss to the project, and the new instrument now reports both halves without having to choose between them. It also exposes a limit of the method, honestly stated: a shared roof deck is easy to cost and hard to capitalize, and the qualitative score can carry the intention only so far before the area, cost, and coverage lines reassert themselves.

Move forward The first four schemes have tested expression and program at modest height. Iteration 5 tests the last lever the architect has not pulled: scale. It carries the full formal language upward to twenty-four stories and adds roughly thirty-six units, asking whether spreading the architecture over a much larger leasable base finally amortizes it, or whether height brings its own costs faster than it brings area.

5.5 Iteration 5 – Height Extrusion

Architectural move Iteration 5 extrudes the parti to twenty-four stories, roughly 278 feet, carrying the swept corner, the balcony expression, and the stepped massing logic upward at full strength. This is the most ambitious scheme in the set and the largest, at 160 units and 149,617 sf of leasable area. It is not a regularization or a retreat to efficient floors; it is the architectural language taken as far up as the assumed envelope will allow. The design question is whether height, which adds both leasable area and prominence, can finally make the accumulated expression pay.




Figure 23: Tall Order

Geometry

Stories	24
Approx. height	278 ft
Setbacks	10' front/back, 15' sides
Massing	Iteration 4 extruded to 24 stories
Site area	16,078 sf
Gross floor area	241,272 sf
Leasable floor area	149,617 sf
Typical plate	≈ 10,053 sf/floor
Typical leasable area	≈ 6,234 sf/floor
Core	≈ 946 sf/floor
Units	160




Figure 24: Aerial view, Iteration 5.



Figure 25: Floor plan, Iteration 5.

Index extraction $C = 0.750$, $Q = 0.536$. Complexity reaches the highest value in the sequence by a wide margin, and the cause is unambiguous: the height premium is now $\chi = 1.917$, nearly doubling a pre-height complexity of 0.392 to a final $C = 0.750$. Height also amplifies the geometric submetrics, lifting stack variation to 0.156 and edge variation to 0.040 as the tall stack accumulates plate-to-plate change. Quality reaches its peak as well, $Q = 0.536$, driven almost entirely by the rate side: the solar score saturates near unity for the tallest scheme and the qualitative score holds at 0.667, pushing Q_R to 0.832. The area side remains the laggard at $Q_A = 0.396$, held down by the lowest average leasable area per floor in the set. The result is the most expressive and highest-quality building in the study, carried on the highest complexity the framework has measured.

Iteration 5 – Height Extrusion	
<i>Architectural Complexity</i>	Value
Envelope, C_{env}	0.045
Stack variation, C_{stack}	0.156
Edge variation, C_{edge}	0.040
Core offset, C_{core}	0.073
Height multiplier, χ	1.917
Pre-height geometric complexity, C_G	0.392
Composite Architectural Complexity, C	0.750
<i>Asset Quality</i>	Value
Area yield, $\hat{Q}_{\text{yield}}$	0.388
Area efficiency, $\hat{Q}_{\text{eff}}$	0.405
Solar exposure, Q_{sun}	0.521
Qualitative, Q_{qual}	0.667
Area-side composite, Q_A	0.396
Rate-side composite, Q_R	0.832
Composite Asset Quality, Q	0.536

Diagnostic readout

Iteration 5 – Height Extrusion	24 stories
Median IRR	-5.2%
P10 / P90 IRR	-22.4% / 4.5%
σ_{IRR}	9.4%
$P(\text{IRR} \geq \tau)$	1%
Median DSCR	0.79×
$P(\text{Robust} \mid C, Q)$	0%
Yield-on-Cost (YOC)	4.08%
Total Development Cost (TDC)	\$175.9M
NOI (stabilized)	\$7.19M
Efficacy, $\mathcal{E}$ ($\mathcal{E}_{\text{rel}}$)	0.220 (0.718)




Figure 26: IRR distribution for Iteration 5, baseline shown for reference (grey). The entire distribution has shifted below the hurdle, with a median return in negative territory.

Reading Height is where ambition stops being amortizable. The largest scheme is also the worst performer: median IRR is now negative at -5.2% , almost no draws clear the hurdle ($P(\text{IRR} \geq \tau) = 1\%$), robustness is zero, and median DSCR has fallen to $0.79\times$. Efficacy reaches its lowest point at $\mathcal{E}_{\text{rel}} = 0.718$. The diagnostic is not subtle, and neither is the cause. Both leasable area and NOI are the highest in the sequence, at 149,617 sf and \$7.19M, yet total development cost has climbed to \$175.9M, an order of magnitude past the income's ability to service it. The height premium that doubled the complexity index has a direct financial twin: the cost of building to twenty-four stories on a 16,000 sf lot rises faster than the rent the added floors can produce, so yield-on-cost falls to 4.08% even as gross income rises. This is the height elasticity of cost made visible. The added area is real, the added prominence is real, and neither is enough, because the capital required to reach them grows faster than both. The model does not reject height in principle. It reports, in this geometry and at this efficiency, that height is the most expensive way the project has found to lose money.

Move forward Two questions remain, and the chapter takes them in order. The first is positional: where do these five schemes actually fall on the efficacy surface, and what does the shape of their trajectory say that the individual readouts do not? The Trajectory Summary takes that up. The second is more fundamental. Every scheme above has been underwritten as a rental asset against a $1.25\times$ debt-service floor, and that floor has been quietly fatal from the baseline onward. The closing section asks whether the tenure assumption is part of the result by re-running all five schemes as for-sale condominiums.

5.6 Trajectory Summary

Read end to end, the sequence tells a cleaner story than any single iteration. The pro forma declines monotonically. Median IRR falls at every step, from 11.7% to -5.2%; robustness falls at every step, from 35% to zero; median DSCR falls at every step, from $1.22\times$ to $0.79\times$. Quality, meanwhile, rises almost monotonically, from 0.452 to 0.536, because each architectural move was authored to improve the building and each one did. The experiment is therefore not a search for an optimum. It is a record of architectural decisions being priced, and the headline is that more architecture bought more quality and less feasibility at nearly every turn.

The efficacy index is the through-line that keeps that headline from being a caricature. Plotted across the sequence, $\mathcal{E}_{\text{rel}}$ does not fall monotonically. It rises first, to 1.030 at the balconies and 1.029 at the swept corner, then falls, to 0.916 at the amenity deck and 0.718 at the height extrusion. The two cheap articulation moves are efficacy-positive even though they are cash-flow-negative, because they buy quality with almost no complexity and lose only leasable area, a revenue effect the efficacy surface does not see. The two expensive structural moves are efficacy-negative, because they raise complexity and capital faster than they raise quality. The framework thus separates two kinds of architectural ambition that a single feasibility number would have flattened together: the kind that is affordable in complexity terms and merely costs revenue, and the kind that is expensive in complexity, capital, and revenue at once.

Table 11: Diagnostic readout across the five design iterations (apartment tenure).

Iter.	C	Q	$\mathcal{E}$	$\mathcal{E}_{\text{rel}}$	IRR	DSCR	$P(R)$	Move
1	0.162	0.452	0.307	1.000	11.7%	1.22	35%	Baseline
2	0.165	0.475	0.316	1.030	7.0%	1.08	11%	Perimeter balconies
3	0.155	0.471	0.316	1.029	6.3%	1.06	9%	Swept southwest corner
4	0.387	0.495	0.281	0.916	1.6%	0.96	2%	Rooftop amenity deck
5	0.750	0.536	0.220	0.718	-5.2%	0.79	0%	Height extrusion

Table 12: Change between sequential iterations, reported as differences of point estimates (apartment tenure). These are deltas of the summary statistics in Table 11, not paired common-random-number transition diagnostics.

Transition	ΔC	ΔQ	$\Delta \text{NLA (sf)}$	$\Delta \text{med. IRR}$	$\Delta P(R)$
Iteration 1 $\rightarrow$ 2	+0.003	+0.023	-13,882	-4.7 pp	-24 pp
Iteration 2 $\rightarrow$ 3	-0.010	-0.004	-3,394	-0.7 pp	-2 pp
Iteration 3 $\rightarrow$ 4	+0.232	+0.024	+6,036	-4.7 pp	-7 pp
Iteration 4 $\rightarrow$ 5	+0.363	+0.041	+39,807	-6.8 pp	-2 pp

The change table makes the geometry of the trajectory explicit. The first two moves cost almost nothing in complexity (a ΔC of $+0.003$ and -0.010) and almost everything in area: the balcony step alone sheds nearly 14,000 sf of leasable floor and 24 points of robustness. The last two moves invert that profile. The amenity deck and the height extrusion add leasable area, substantially so in the case of height, but they do it by spending complexity at a rate the earlier moves never approached, $+0.232$ and $+0.363$, with the height multiplier doing most of the damage. The robustness column is nearly exhausted by Iteration 3, which is why its later decrements look small: there was little left to lose. Read together with the comparative distributions in Figure 27, the picture is of a set of returns that not only shift left at every step but widen as they go, with σ_{IRR} rising from 7.7% at the baseline to 10.0% at the amenity deck, so that the more ambitious schemes are both lower and less certain.




Figure 27: Comparative IRR distributions across the five iterations (apartment tenure): cumulative (left) and density (right). The amenity deck is bimodal because the discrete delay and absorption draws split the runs into two clusters that straddle breakeven at its thin margin; the height extrusion’s mode shifts negative and its left tail fattens because the multiplicative cost shock compounds against the height-cost multiplier, widening the absolute spread on a taller, costlier scheme.

The trajectory plotted on the efficacy surface is the chapter’s closing act for the rental sequence: five architectural decisions priced by a stochastic financial model and located on the instrument built in Chapter 3. The schemes do not scatter across the surface; they trace a path. The articulation phase moves the project up the quality axis while complexity holds within a hundredth of the baseline, and the structural phase then turns hard to the right, driving complexity from 0.155 to 0.750 while quality inches from 0.471 to 0.536. The iso-efficacy contours give that turn a price, which is the level-set reading promised in §3.9. To hold the swept corner’s efficacy at the amenity deck’s complexity of $C = 0.387$, the surface requires a quality of roughly 0.60; the scheme delivers 0.495, about a quarter of the required gain. To hold even baseline efficacy at the height extrusion’s $C = 0.750$, quality would need to reach approximately 0.92, near the ceiling of the index; the scheme

posts 0.536.




Figure 28: Safe Design Envelope, efficacy surface with the five schemes plotted in (C, Q) space (apartment tenure).

The surface also reports on the calibration itself, and this is the quieter finding. With the market constants pinned to the Eastlake baseline, the maximum attainable efficacy anywhere on the surface is $\Lambda_R/\Lambda_K \approx 0.53$, reached only at the unbuildable corner of zero complexity and perfect quality. The break-even contour $\mathcal{E} = 1$, where quality-lifted yield balances complexity-loaded cost, lies entirely off the map. No admissible geometry reaches it, which is the surface-level statement of what every diagnostic in this chapter has reported pointwise: under rental tenure, on this site, at this cost basis, the design problem is being asked to solve an underwriting problem it cannot reach. The same conclusion arrives through the robustness lens, since the best of the five schemes clears only 35% of sampled futures and none enters an envelope drawn at any conventional threshold. The one move in the study that lifts a scheme inside such an envelope is not geometric at all, which is the question the next section takes up.

It should be said plainly that this is not the population the surface was built for. The instrument of Chapter 3 anticipates a Mode A sweep across the full (C, Q) grid and a Mode B record dense enough to trace the envelope boundary empirically: several distinct partis, each carried through many iterations, with massing generation automated in Grasshopper so that authored schemes cross the boundary from multiple directions. This experiment supplies one parti family and five points, three of them within $\Delta C = 0.01$ of one another, which is a path rather than a population. Within-project, the diagnostics operate at full resolution, and the readings above are exact for these schemes. What five points along one path cannot do is locate the boundary, distinguish the behavior of move types, or support inference beyond this site, the restriction §3.11 placed on the analysis at the outset. The figure should therefore be read as a demonstration that the measurement infrastructure works at the resolution the methods chapter claimed, with the populated surface standing as the experiment this one was scoped down from, and the obvious next run.

5.7 Tenure Test: Apartment Baseline Against Condominiums

The rental sequence carries one assumption that never changed and never stopped mattering: every scheme was underwritten as an apartment against a $1.25\times$ debt-service floor. To test whether that assumption is part of the result rather than a neutral backdrop, all five schemes were re-run as for-sale condominiums, holding the geometry, the cost basis, and the stochastic inputs fixed and changing only the tenure logic. For the for-sale variant the income line becomes net sale proceeds rather than stabilized NOI, the coverage test becomes a developer profit margin rather than DSCR, and the hold shortens to a development-and-sellout period. The efficacy surface itself does not move: efficacy is a function of C , Q , and the baseline economic calibration, so $\mathcal{E}_{\text{rel}}$ is the same 1.000, 1.030, 1.029, 0.916, 0.718 under either tenure. What changes is where the robustness probability lands on top of it.

At the baseline the change is dramatic and runs in the project's favor. The condominium baseline returns a median IRR of 14.9% against the apartment's 11.7%, a developer margin of 23%, and a robustness of 64% against the apartment's 35%. It is the strongest single result in the entire study, and for once it clears the hurdle comfortably. The reason is the constraint that was quietly fatal to the rental sequence. Under apartment tenure the baseline saw 49% of draws clear the IRR hurdle but only 35% pass jointly, the fourteen-point gap being the DSCR floor cutting into otherwise acceptable returns. Under condominium tenure that gap nearly vanishes: 64% of draws clear the hurdle and 64% are robust, because the for-sale structure carries no debt-service floor and the margin test rarely binds when the return is healthy. For-sale tenure does for the disciplined baseline what no amount of articulation could. It removes the drag that kept the rental baseline marginal, and it does so by capturing value at sale rather than amortizing it across a leveraged hold (Figure 29).




Figure 29: Baseline IRR distribution by tenure: apartment (rental) against condominium (for-sale), cumulative (left) and density (right). The condominium distribution sits well to the right of the rental one, with more mass above the hurdle.

The favor does not survive the architectural moves. If anything, the for-sale model is less forgiving of them than the rental model, and the reason is structural. A rental pro forma amortizes a loss of leasable area across a long hold and a terminal value; a for-sale pro forma has to recover every dollar at sale, so each square foot of sellable area removed by an architectural gesture is felt immediately and in full. The result is a steeper decline. Where the rental schemes drift down, the condominium schemes fall off a ledge. The balcony and swept-corner schemes still earn positive but thin margins of 9% and 7%. The amenity deck turns the project negative, with a developer margin of -3% and net sale proceeds of roughly $-\$0.9\text{M}$, meaning the building would be sold for less than it cost to build. The height extrusion is worse still, at a -21% margin and roughly $-\$10.2\text{M}$ in net proceeds, a catastrophic outcome that the rental version, buffered by its hold, registered only as a deeply negative return. Table 13 sets the two tenures side by side.

Table 13: Tenure comparison across the five iterations. Apartment robustness is governed by the $1.25\times$ DSCR floor; condominium robustness is governed by a developer-margin floor. Geometry, cost basis, and stochastic inputs are identical across tenures.

Iter.	Move	Apartment (rental)		Condominium (for-sale)		
		IRR	$P(R)$	IRR	Margin	$P(R)$
1	Trapezoidal baseline	11.7%	35%	14.9%	23%	64%
2	Perimeter balconies	7.0%	11%	6.5%	9%	26%
3	Swept southwest corner	6.3%	9%	5.3%	7%	22%
4	Rooftop amenity deck	1.6%	2%	-2.4%	-3%	5%
5	Height extrusion	-5.2%	0%	-21.9%	-21%	0%

One asymmetry the comparison cannot register deserves explicit notice, and practitioner review of this experiment raised it directly: renters and buyers do not price the same architecture identically. The model is tenure-aware only at the cap, with revenue premiums of 15% on rents and 20% on sale prices (\$4.5), while the indices themselves are tenure-blind; a wraparound balcony scores the same Q whether a tenant rents it or a buyer owns it. The for-sale literature suggests the difference runs deeper than a cap. Buyers capitalize vertical position, corner exposure, and floor level into transaction prices at premiums rental markets do not reproduce [Conroy et al., 2013, Wong et al., 2011, Ben-Shahar et al., 2023], daylight is priced into housing sales independent of size and location [Loro et al., 2024], and the rental design premiums documented for income property are real but thinner [Nase et al., 2016].

One developer reviewing this work put it concretely: condominium buyers will pay for generous private outdoor space in a way tenants rarely do, because the balcony is purchased rather than borrowed. If that is right, the condominium readings for the balcony and swept-corner schemes are conservative. Those are precisely the schemes whose lost sellable area the model prices in full while crediting the offsetting buyer premium only up to the capped, tenure-blind uplift, and modest re-weighting moves them: price premiums of roughly 5 to 8% beyond the current credit would carry both schemes to the 15% margin floor. The qualification has a hard edge, though. The amenity deck would need nearly triple the modeled premium to reach the same floor, and the height extrusion would need a 27% price uplift merely to break even, so no plausible buyer enthusiasm rescues the structural moves. The middle of Table 13 should therefore be read as a floor rather than a verdict, and the gap it exposes names a concrete calibration refinement: tenure-specific premium functions, estimated separately from sale and rent hedonics, sitting on the existing Q_A/Q_R channels without any change to the framework's structure.

The sensitivity structure confirms that the two tenures are exposed to different risks even on the same building. For the apartment baseline the widest swings come from rent growth, which moves median IRR from roughly -4.7% to $+25.8\%$ across its range, followed by hard cost and base rent. For the condominium baseline the dominant lever is market drift, the sale-price appreciation assumption, which moves median IRR from about 5.9% to 24.6% , followed by hard cost and sale price. The common term is hard cost: at \$420 per sf it is among the two or three widest bars under either tenure, which is the quantitative reason both tenures punish the height extrusion so severely. The rental model worries most about whether rents will grow into the debt; the for-sale model worries most about whether the market will hold through sellout. Both worry about what it costs to build, and on a tall, slender tower they have every reason to.

The tenure test therefore closes the chapter on a result that neither sequence could produce alone. The decision of how to hold the asset is not separable from the decision of how to shape it. For a disciplined, area-efficient baseline, switching to for-sale is the single most powerful feasibility move in the study, more powerful than any geometric change, because it lifts the debt-service constraint that the rental structure imposes from the start. For an ambitious, area-sacrificing scheme, the

same switch is a liability, because the for-sale model gives back nothing at the hold to cushion the lost area. The efficacy surface stays fixed under both tenures, as it should, since the geometry's conversion of complexity into quality has nothing to do with how the building is sold. The robustness map slides across that fixed surface as the tenure changes. That separation, a stable efficacy surface beneath a tenure-dependent robustness field, is the most compact statement of what the two-mode framework is for: it lets the architect and the developer see, at the same time and on the same building, which problems are problems of form and which are problems of finance.

6 Conclusion: Flying the Envelope

This thesis began with a narrow complaint that turned into a method. The American development pro forma does not really see architecture. It sees gross floor area, story count, cost per square foot, rent per square foot, cap rate, loan terms, and exit value [Pyhrr, 1973, Miles et al., 2007]. Those numbers are real, but they are the residue a building leaves once its geometry has already decided most of what the numbers will say. The parti settles massing, footprint, core logic, structural attitude, envelope burden, daylight, and vertical ambition long before the project is detailed enough to underwrite [Ward et al., 1995, Sobek et al., 1999]. The tools available at that moment are either architectural and qualitative or financial and scalar, and the two rarely meet. The literature already accepts that form is not financially neutral, that markets price spatial and formal attributes, only usually after the building is finished and transacted [Rosen, 1974, Rong et al., 2020, Nase et al., 2016].

The Safe Design Envelope tries to move that pricing earlier. It encodes the parti as two design-state variables, Architectural Complexity, C , and Asset Quality, Q . Complexity carries the cost-facing consequences of geometry: envelope burden, stack and edge variation, and core eccentricity [Belniak et al., 2013, Eriksen and Orlando, 2021]. Quality carries the value-facing consequences: yield, efficiency, daylight, and authored intentions such as green ambition and material specificity [Loro et al., 2024, Kahn and Kok, 2014, Nase et al., 2016]. The indices do not capture architecture in full and are not meant to. They capture enough of a building's cost- and value-relevant structure to let a design enter a probabilistic pro forma while it is still a design.

This places the work between two settled literatures rather than above them. Probabilistic real estate modeling has long held that feasibility is a distribution and not a verdict [Pyhrr, 1973, Hoesli et al., 2006, Loizou and French, 2012, Leung, 2014, French and Gabrielli, 2006]. Hedonic and design-valuation research has long held that form carries price [Rosen, 1974, Rong et al., 2020, Nase et al., 2016]. Leung [2014] made the pro forma probabilistic but left the building scalar; Rong et al. [2020] made form financially visible but did so after the fact. The contribution here is the seam between them, and it is a small seam: a prospective, geometry-aware, uncertainty-based instrument for schematic design, assembled from methods that are not themselves new [Geltner and de Neufville, 2018, Dixit and Pindyck, 1994].

It is not an optimizer and not a machine for finding the one correct building. It is a bounded region of design space within which ambition can be pursued with some quantified awareness of cost, revenue, coverage, and market uncertainty. It does not replace judgment. It gives judgment a few more instruments to read, and it should be received as such.

6.1 What the Demonstrator Actually Showed

The experiment was modest, and not only by intention. Time and tooling limited it to a single family of moves; a fuller study would evaluate three or four distinct partis

across more iterations, which is achievable with more automation in the Grasshopper pipeline. As run, it used one constrained Eastlake site, one hypothetical zoning envelope, one calibrated pro forma, and one Monte Carlo framework. The five schemes were drawn as architectural propositions first and evaluated through identical financial machinery second. That order matters: the experiment was not a search for the most profitable shape but a test of whether the financial consequences of authorship could be made visible while the design was still legible as design.

The first result is blunt. Under apartment tenure, the sequence grew less financially robust as it grew more architecturally *interesting*. The trapezoidal baseline returned a median IRR of 11.7% and $P(\text{Robust}) = 35\%$. Balconies improved spatial quality and lowered robustness to 11%; the swept corner held formal interest and fell to 9%; the rooftop amenity reached 2%; the height extrusion collapsed, with a median IRR of -5.2% , a median DSCR of $0.79\times$, and $P(\text{Robust}) = 0\%$. The model found no hidden path through the pro forma. It found the wall.

That negative result is among the more useful ones. A framework that only flatters design moves is not an analytical instrument. The envelope showed that balconies, rounded corners, amenity decks, and added height do not become defensible simply because they improve quality. They have to add value faster than they add cost, remove area, or strain coverage, and here most did not. Height failed most clearly. Iteration 5 raised leasable area and stabilized NOI, but capital rose faster, the familiar height-cost elasticity documented in skyscraper economics, observed here at the scale of a single parti [Ahlfeldt and Barr, 2022, Ahlfeldt and McMillen, 2018, Flanagan and Norman, 1978, Blackman and Picken, 2010]. On this site and under this cost structure, height was the most expensive way to lose money.

The second result is quieter. The Efficacy surface and the robustness field did not always agree. Some moves improved the deterministic quality-to-complexity relationship while still weakening probabilistic feasibility. That divergence is the point. Efficacy asks whether a move is efficient in design-value terms; robustness, formally a chance constraint on the joint return and coverage thresholds, asks whether the move survives stochastic cost, rent, financing, and absorption [Charnes and Cooper, 1959, Prékopa, 1994, Geltner and de Neufville, 2018]. When the two separate, the architect learns something exact: a move can be spatially promising and financially fragile at once. The framework does not resolve that tension. It names it early.

The third result came from tenure. Re-run as for-sale condominiums, the baseline improved sharply, to a median IRR of 14.9%, a developer margin of 23%, and $P(\text{Robust}) = 64\%$. This was the single strongest feasibility move in the study, and it was not a geometric move at all. It was disposition and capital structure, consistent with the premium that vertical position and floor level command in for-sale markets [Conroy et al., 2013, Wong et al., 2011, Ben-Shahar et al., 2023]. The advantage did not survive the area-sacrificing schemes: the amenity deck turned negative and the height extrusion turned catastrophic. Form and tenure are not separable. The same building can be marginal as a rental, strong as a condominium, or unbuildable as either once the geometry spends sellable area too freely. The sharpest risks sat at the interface, where each domain assumed the other would

absorb the loss.

6.2 What This Thesis Contributes

The thesis contributes three modest things.

First, a method for translating early form into financially legible variables without collapsing design into taste. The C and Q system is a measurement scaffold for cost- and value-facing consequences, assembled from established composite-indicator and multi-attribute practice rather than invented whole [OECD and European Commission, Joint Research Centre, 2008, Keeney and Raiffa, 1976], and it joins a small body of work that treats architectural form as something quantifiable [Pakkanen, 2013, Yoo et al., 2025]. It concedes that architecture holds qualities no model can capture, while refusing the opposite error of treating geometry as financially mute [Rosen, 1974, Nase et al., 2016].

Second, it treats feasibility as a region rather than a verdict. A conventional pro forma returns a binary: the project pencils or it does not. The envelope instead maps the probability that a design state clears both the return and the coverage hurdle across many market futures, following set-based design and tradespace exploration, where the aim is to defer commitment and preserve a defensible region while uncertainty is high [Ward et al., 1995, Sobek et al., 1999, Ross and Hastings, 2005, Mourtzis et al., 2020, Østergård et al., 2015]. In architectural terms, it keeps the parti soft a little longer.

Third, it adjusts the architect's role in the financial conversation. The architect is not asked to become a lender or an appraiser, only to understand how formal decisions propagate through cost, revenue, and risk. That is a kind of technical co-authorship rather than refusal from outside the room or compliance after the room has emptied. Ignoring finance does not protect architecture; it only lets architecture enter the model as an unpriced liability. Submitting entirely to deterministic underwriting is not realism but surrender. The envelope proposes a third stance: authorship under uncertainty, with risk read as one of the media through which form is now composed.

6.3 Limitations of the First Pass

The framework is a first pass, and its limits set the agenda.

It uses one site. Eastlake is useful because it is constrained, urban, and height-sensitive, but no single parcel validates a general claim, and the method needs testing across land costs, zoning envelopes, submarkets, and construction types, since land price and regulation drive feasibility as forcefully as form does [Mayer and Somerville, 2000, Huang, 2021, Titman, 1985]. The zoning envelope is hypothetical, chosen to test financial behavior rather than entitlement, so the result is not a development proposal. The C and Q indices still need empirical calibration against real cost estimates, facade takeoffs, structural quantities, and lease-up and transaction data; their weights and normalization remain experimental [OECD and European Commission, Joint Research Centre, 2008, RSMeans, 2024, Leśniak and

Górka, 2021], and the authored quality scores smuggle judgment into the model until they are supported by rubrics, inter-rater testing, or survey data [Keeney and Raiffa, 1976, Barron and Barrett, 1996, Stillwell et al., 1981]. The stochastic model simplifies correlation: cost, rates, rent, absorption, and cap rates move together, and the dependency structure should be widened through correlated draws or regime sampling [Touran and Wiser, 1992, Marino et al., 2008]. The pro forma itself stays schematic, a decision aid for when the design is still changing rather than a substitute for a budget, schedule, or term sheet [Miles et al., 2007, Moorhead et al., 2024].

The deepest limitation is not technical. This thesis quantifies architecture for economic gain and does not examine what that means. It accepts the developer's scoreboard, IRR, DSCR, and margin, and asks only how form scores on it. It does not engage the politics of measurement, the housing-affordability stakes riding on which buildings get financed, or the plain question of whether making architecture legible to capital expands the discipline's agency or actually makes architecture a more obedient input [Dong, 2023, Mayer and Somerville, 2000, United Nations Development Programme, 2010]. Those are the political questions the premise raises, and they are bracketed here. The bracketing is the honest boundary of the work, and probably the more important thesis standing behind this one.

6.4 Disciplinary Agency After the Pro Forma

The broader implication is disciplinary, and it is better stated plainly than dressed up. Architecture tends to take one of two weak positions in speculative development. One is heroic refusal, which makes an aesthetic claim and ignores the machinery deciding whether the building will exist. The other is quiet compliance, which manages surfaces after the financial model has already narrowed the field. Both give up too much.

The pro forma is already shaping the building, so it is more usefully treated as a design medium than as an external verdict [Geltner and de Neufville, 2018, Leung, 2014]. The envelope lets the architect argue upstream in financial terms: this massing changes envelope burden by this much; this balcony raises quality but removes revenue area; this height move lifts NOI and destroys yield-on-cost; this tenure shift rescues the baseline but punishes lost area. That is a harder argument to wave off than a claim of better taste.

It is worth being blunt about the cost of this move. Teaching architecture to speak finance does not free architecture from finance. It may only make the architect more useful to the people who were already going to decide whether the building exists. Whether that is leverage or capture is not a question a spreadsheet can settle, and this thesis does not settle it. It builds the instrument and leaves the politics open. That is a small contribution offered to a much larger argument the discipline has barely begun.

6.5 The Demonstrator

The aerospace metaphor that opened the thesis is still useful, if held loosely. A demonstrator is not a product. It shows that a category can be entered again, and that an old verdict may have mistaken a historical failure for a physical limit [Boom Supersonic, 2025].

The Safe Design Envelope is that and no more. It does not reform construction costs, stabilize interest rates, or fix the zoning politics, entitlement delay, and risk aversion that make ambitious urban housing hard to deliver in American markets [Mayer and Somerville, 2000, Dong, 2023]. Those conditions remain political, and no model substitutes for the work of changing them.

What the thesis shows is smaller and still worth saying. The conflict between architectural ambition and financial feasibility is not unknowable. It can be measured early, mapped, and decomposed into form, market, tenure, and financing problems. It can show when quality is creating value, when complexity is destroying it, and when a project fails not because of its architecture but because the deal has asked the architecture to carry a load it cannot bear.

That is the limited optimism here. Not that better architecture will always pencil, because it will not. Not that simulation can save the city, because it cannot. Only that the discipline can build better instruments for a fight it is already in, and that it should not mistake a quiet instrument for a loud achievement. The prototype flew, found the edge of the envelope sooner than expected, and put one scheme straight into the ground. That is what prototypes are for. The next version should be broader, better calibrated, less forgiving, and braver about the politics this one set aside.

References

- G. M. Ahlfeldt and J. Barr. The economics of skyscrapers: A synthesis. *Journal of Urban Economics*, 129:103419, 2022. doi: 10.1016/j.jue.2021.103419. URL <https://doi.org/10.1016/j.jue.2021.103419>.
- G. M. Ahlfeldt and D. P. McMillen. Tall buildings and land values: Height and construction cost elasticities in chicago, 1870–2010. *Review of Economics and Statistics*, 100(5):861–875, 2018. doi: 10.1162/rest.a.00734. URL https://doi.org/10.1162/rest_a_00734.
- M. M. Ali and K. S. Moon. Structural developments in tall buildings: Current trends and future prospects. *Architectural Science Review*, 50(3):205–223, 2007. doi: 10.3763/asre.2007.5027. URL <https://doi.org/10.3763/asre.2007.5027>.
- T. A. T. Awida. The effect of core eccentricity on the structural behavior of concrete tall buildings. *Journal of Civil Engineering and Architecture*, 18(3):112–117, 2024. doi: 10.17265/1934-7359/2024.03.002. URL <https://doi.org/10.17265/1934-7359/2024.03.002>.
- F. H. Barron and B. E. Barrett. Decision quality using ranked attribute weights. *Management Science*, 42(11):1515–1523, 1996. doi: 10.1287/mnsc.42.11.1515. URL <https://doi.org/10.1287/mnsc.42.11.1515>.
- S. Belniak, A. Leśniak, E. Plebankiewicz, and K. Zima. The influence of the building shape on the costs of its construction. *Journal of Financial Management of Property and Construction*, 18(1), 2013. doi: 10.1108/13664381311305096.
- D. Ben-Shahar, Y. Deng, and E. Sulganik. Vertical status in real estate. *Real Estate Economics*, 51(3):747–773, may 2023. doi: 10.1111/1540-6229.12422.
- I. Q. Blackman and D. H. Picken. Height and construction costs of residential high-rise buildings in Shanghai. *Journal of Construction Engineering and Management*, 136(11):1169–1180, 2010. doi: 10.1061/(ASCE)CO.1943-7862.0000226. URL [https://doi.org/10.1061/\(ASCE\)CO.1943-7862.0000226](https://doi.org/10.1061/(ASCE)CO.1943-7862.0000226).
- Boom Supersonic. XB-1 achieves supersonic flight. Boom Supersonic press release, 28 January, 2025. First supersonic flight, Mojave Air & Space Port; reached Mach 1.12 at 35,290 ft. Boomless Cruise (no boom reaching ground) announced 10 February 2025.
- W. B. Brueggeman and J. D. Fisher. *Real Estate Finance and Investments*. McGraw-Hill Education, New York, NY, 17th edition, 2021.
- Building Owners and Managers Association International. Boma standards. <https://www.boma.org/boma-standards/>, n.d. Accessed 2026-06-07.

- S. Burhenne, D. Jacob, and G. P. Henze. Uncertainty analysis in building simulation with monte carlo techniques. In *Proceedings of SimBuild 2010, Fourth National Conference of IBPSA-USA*, 2010.
- CBRE Research. Pacific northwest multifamily market outlook. Technical report, CBRE, 2024. Grey literature; verify exact title and year.
- M. A. V. Chaban. Flatiron Building, Admired but Rarely Copied, Inspires Developers. *The New York Times*, April 4 2016. URL <https://www.nytimes.com/2016/04/05/nyregion/flatiron-building-admired-but-rarely-copied-inspires-developers.html>.
- A. Charnes and W. W. Cooper. Chance-constrained programming. *Management Science*, 6(1):73–79, 1959. doi: <https://www.jstor.org/stable/3088446>.
- K. W. Chau, S. K. Wong, Y. Yau, and A. K. C. Yeung. Determining optimal building height. *Urban Studies*, 44(3):591–607, 2007. doi: 10.1080/00420980601131902.
- J. Clayton, D. C. Ling, and A. Naranjo. Commercial real estate valuation: Fundamentals versus investor sentiment. *Journal of Real Estate Finance and Economics*, 38(1):5–37, 2009. doi: 10.2139/ssrn.1132361.
- S. J. Conroy, A. Narwold, and J. Sandy. The value of a floor: Valuing floor level in high-rise condominiums in san diego. *International Journal of Housing Markets and Analysis*, 6(2):197–208, 2013. doi: 10.1108/IJHMA-01-2012-0003.
- Cushman & Wakefield. Seattle Multifamily MarketBeat Q4 2025. Cushman & Wakefield Research, 2025. URL https://assets.cushmanwakefield.com/-/media/cw/marketbeat-pdfs/2025/q4/us-reports/multifamily/seattle_americas_marketbeat_multifamily_q42025.pdf. Public multifamily market report for the Seattle region, including rent, vacancy, construction pipeline, and employment context.
- B. D’Amico and F. Pomponi. A compactness measure of sustainable building forms. *Royal Society Open Science*, 6(6):181265, 2019. doi: 10.1098/rsos.181265. URL <https://doi.org/10.1098/rsos.181265>.
- P. de Jong and S. C. F. van Oss. High rise costs. In *Proceedings of the 4th International SCRI Symposium*, Salford, UK, 2007. University of Salford. Also archived via CIB, irbnet CIB16569.
- A. K. Dixit and R. S. Pindyck. *Investment Under Uncertainty*. Princeton University Press, Princeton, NJ, 1994. doi: <https://doi.org/10.2307/j.ctt7sncv>.
- DOCOMOMO US. Streamline Moderne. DOCOMOMO US Style Guide, 2026. URL <https://www.docomomo-us.org/style/>

- streamline-moderne. Reference for Streamline Moderne as a horizontal style defined by curved surfaces, curved corners, cylindrical projections, and restrained ornament.
- H. Dong. Higher-density development for lower-cost housing? Understanding the multifamily housing market and the role of density in multifamily home prices. *Journal of Planning Education and Research*, 43(3):617–636, sep 2023. doi: 10.1177/0739456X20912829.
- S. F. Eletter et al. Using monte carlo simulation for real estate price valuation. In *Proceedings of the 25th International Arab Conference on Information Technology (ACIT)*. IEEE, 2024. doi: 10.1109/ACIT62805.2024.10876955. Verify full author list and page range.
- M. D. Eriksen and A. W. Orlando. Returns to scale in residential construction: The marginal impact of building height. *Real Estate Economics*, 49(4):1135–1164, 2021. doi: <https://doi.org/10.1111/1540-6229.12357>.
- European Commission, Joint Research Centre. Step 7: Aggregating indicators. https://knowledge4policy.ec.europa.eu/composite-indicators/toolkit_en/navigation-page/10-step-guide_en/step-7-aggregating-indicators_en, n.d. Competence Centre on Composite Indicators and Scoreboards. Accessed 2026-06-07.
- Fannie Mae. Conventional properties term sheet. Online term sheet, 2026. URL <https://multifamily.fanniemae.com>. Accessed 2026. Verify access date and specific URL before final submission.
- R. Flanagan and G. Norman. The relationship between construction price and height. *Chartered Surveyor Building and Quantity Surveying Quarterly*, 54:68–71, 1978.
- Freddie Mac. Small balance loan term sheet. Online term sheet, 2026. URL <https://mf.freddiemac.com>. Accessed 2026. Verify access date and specific URL before final submission.
- N. French and L. Gabrielli. Uncertainty and feasibility studies: An italian case study. *Journal of Property Investment & Finance*, 24(1):49–67, 2006. doi: 10.1108/14635780610700732. URL <https://doi.org/10.1108/14635780610700732>.
- D. Geltner and R. de Neufville. *Flexibility and Real Estate Valuation under Uncertainty: A Practical Guide for Developers*. Wiley-Blackwell, 2018. doi: 10.1002/9781119106470. URL <https://doi.org/10.1002/9781119106470>.
- D. Geltner, N. G. Miller, J. Clayton, and P. Eichholtz. *Commercial Real Estate Analysis and Investments*. OnCourse Learning, Mason, OH, 3rd edition, 2014.

- D. Gimpelevich. Simulation-based excess return model for real estate development. *Journal of Property Investment & Finance*, 29(2):115–144, 2011. doi: 10.1108/14635781111112765. URL <https://doi.org/10.1108/14635781111112765>.
- P. Glasserman and D. D. Yao. Some guidelines and guarantees for common random numbers. *Management Science*, 38(6):884–908, 1992. doi: 10.1287/mnsc.38.6.884. URL <https://doi.org/10.1287/mnsc.38.6.884>.
- C. H. Grohmann, M. J. Smith, and C. Riccomini. Multiscale analysis of topographic surface roughness in the midland valley, scotland. *IEEE Transactions on Geoscience and Remote Sensing*, 49(4):1200–1213, 2011. doi: 10.1109/TGRS.2010.2053546.
- M. Hoesli, E. Jani, and A. Bender. Monte carlo simulations for real estate valuation. *Journal of Property Investment & Finance*, 24(2):102–122, 2006. doi: 10.1108/14635780610655076. URL <https://doi.org/10.1108/14635780610655076>.
- L. Hong et al. Daylight availability of living rooms in dense residential areas under current planning regulations: A cross-region case study in china. *Buildings*, 14(4):1090, 2024. doi: 10.3390/buildings14041090. URL <https://doi.org/10.3390/buildings14041090>.
- Y. Huang. More reliable land price index: Is there a slope effect? *Land*, 10(3):261, 2021. doi: <https://doi.org/10.3390/land10030261>.
- International Code Council. Talking in Code: High-Rise Building Definition. ICC Building Safety Journal, 2019. URL <https://www.iccsafe.org/building-safety-journal/bsj-technical/talking-in-code-high-rise-building-definition/>. Explains the IBC high-rise threshold as an occupied floor more than 75 feet above the lowest level of fire department vehicle access.
- International Code Council. 2021 International Building Code, Chapter 6: Types of Construction. ICC Digital Codes, 2021. URL <https://codes.iccsafe.org/content/IBC2021P2/chapter-6-types-of-construction>. Model-code source for construction-type definitions and fire-resistance logic.
- International Code Council. 2024 International Building Code, Chapter 5: General Building Heights and Areas. ICC Digital Codes, 2024. URL <https://codes.iccsafe.org/content/IBC2024P1/chapter-5-general-building-heights-and-areas>. Model-code source for height, story, and area limits by occupancy and construction type.

- M. E. Kahn and N. Kok. The capitalization of green labels in the california housing market. *Regional Science and Urban Economics*, 47:25–34, 2014. doi: 10.1016/j.regsciurbeco.2013.07.001.
- R. L. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Wiley, New York, 1976.
- S. Kim. Modeling concept for problem-solving, 2020. Cited as presented in Rachmawati and Kim (2022). Verify original publication details before final submission.
- King County GIS Center. King County Parcel Viewer. King County, Washington, 2026. URL <https://kingcounty.gov/en/dept/kcit/data-information-services/gis-center/maps-apps/parcel-viewer>. Official parcel research tool used to locate parcel, assessor, and development-condition information.
- S. Küçükyavuz and R. Jiang. Chance-constrained optimization under limited distributional information: A review of reformulations based on sampling and distributional robustness. *EURO Journal on Computational Optimization*, 10:100030, dec 2022. doi: 10.1016/j.ejco.2022.100030.
- K.-H. Lee and H.-S. Cha. A dynamic simulation model for economic feasibility of apartment development projects. *International Journal of Strategic Property Management*, 23(6):402–414, oct 2019. doi: 10.3846/ijspm.2019.10431.
- A. Leśniak and M. Górka. Pre-design cost modeling of facade systems using the gam method. *Archives of Civil Engineering*, 67(3):123–138, 2021. doi: 10.24425/ace.2021.138047. URL <https://doi.org/10.24425/ace.2021.138047>.
- K. C.-K. Leung. Beyond DCF analysis in real estate financial modeling: Probabilistic evaluation of real estate ventures. S.M. thesis, Massachusetts Institute of Technology, Center for Real Estate, 2014. DSpace handle 1721.1/87612.
- C. Z. Li et al. Schedule risks in prefabrication housing production in hong kong. *Journal of Cleaner Production*, 134:482–494, 2016. doi: 10.1016/j.jclepro.2016.02.123.
- J. Li and Y. Liang. Beyond the cap rate: Valuation of multifamily properties. *Journal of Real Estate Finance and Economics*, 60(1):101–110, 2020. doi: 10.1007/s11146-019-09717-9. URL <https://doi.org/10.1007/s11146-019-09717-9>.
- D. C. Ling and W. R. Archer. *Real Estate Principles: A Value Approach*. McGraw-Hill Education, New York, NY, 6th edition, 2020.
- P. Linneman. *Real Estate Finance and Investments: Risks and Opportunities*. Linneman Associates, Philadelphia, PA, 5th edition, 2018.

- Local Housing Solutions. Gauging development feasibility. <https://www.localhousingsolutions.org/analyze/gauging-development-feasibility/>, n.d. Accessed 2026-06-07.
- P. Loizou and N. French. Risk and uncertainty in development: A critical evaluation of using the monte carlo simulation method as a decision tool in real estate development projects. *Journal of Property Investment & Finance*, 30(2):198–210, 2012. doi: 10.1108/14635781211206922. URL <https://doi.org/10.1108/14635781211206922>.
- S. Loro, V. R. M. Lo Verso, E. Fregonara, and A. Barreca. Influence of daylight on real estate housing prices: A multiple regression model application in turin. *Journal of Building Engineering*, 96:110413, 2024. doi: 10.1016/j.jobbe.2024.110413. URL <https://doi.org/10.1016/j.jobbe.2024.110413>.
- S. Marino, I. B. Hogue, C. J. Ray, and D. E. Kirschner. A methodology for performing global uncertainty and sensitivity analysis in systems biology. *Journal of Theoretical Biology*, 254(2):178–196, 2008. doi: 10.1111/1540-6229.00103.
- C. J. Mayer and C. T. Somerville. Land use regulation and new construction. *Regional Science and Urban Economics*, 30(6):639–662, dec 2000. doi: 10.1016/S0166-0462(00)00055-7.
- M. E. Miles, G. L. Berens, M. J. Eppli, and M. A. Weiss. *Real Estate Development: Principles and Process*. Urban Land Institute, Washington, DC, 4th edition, 2007. Verify author list/spelling against your copy.
- M. Moorhead, L. Armitage, and M. Skitmore. Real estate development feasibility and hurdle rate selection. *Buildings*, 14(4):1045, 2024. doi: 10.3390/buildings14041045. URL <https://doi.org/10.3390/buildings14041045>.
- Mortenson. Seattle Construction Cost Index. Mortenson Cost Index, 2026. URL <https://www.mortenson.com/cost-index/seattle>. Quarterly Seattle construction-cost index used as a public benchmark for local cost escalation conditions.
- G. Mosey and B. Deal. Optimizing multi-family building massing for affordability and envelope performance: An investigation of the trade-offs implicit in low rise residential buildings. *Buildings*, 11(3):99, 2021. doi: 10.3390/buildings11030099. URL <https://doi.org/10.3390/buildings11030099>.
- D. Mourtzis et al. Set-based design: A review and new directions. *Design Science*, 6:e18, 2020. Verify author byline.

- NAIOP. Commercial real estate terms and definitions. <https://www.naiop.org/education-and-career/industry-terms-and-definitions/>, n.d. Entry for “Load Factor or Core Factor.” Accessed 2026-06-07.
- I. Nase, J. Berry, and A. Adair. Impact of quality-led design on real estate value. *Journal of Property Research*, 33(4):309–331, 2016. doi: 10.1080/09599916.2016.1258588. URL <https://doi.org/10.1080/09599916.2016.1258588>.
- National Institute of Standards and Technology. Coefficient of variation. <https://www.itl.nist.gov/div898/software/dataplot/refman2/auxillar/coefvari.htm>, 2017. NIST/SEMATECH e-Handbook of Statistical Methods. Accessed 2026-06-07.
- OECD and European Commission, Joint Research Centre. *Handbook on Constructing Composite Indicators: Methodology and User Guide*. OECD Publishing, Paris, 2008. ISBN 978-92-64-04345-9. doi: 10.1787/9789264043466-en. URL <https://doi.org/10.1787/9789264043466-en>. Primary authors Nardo, Saisana, Saltelli, Tarantola (JRC); Hoffmann, Giovannini (OECD).
- P. Oldfield and B. Doherty. Offset cores: Trends, drivers and frequency in tall buildings. *CTBUH Journal*, (II):40–45, 2019. URL <https://global.ctbuh.org/resources/papers/download/4186-offset-cores-trends-drivers-and-frequency-in-tall-buildings.pdf>.
- T. Østergård, S. E. Maagaard, and R. L. Jensen. A stochastic and holistic method to support decision-making in early building design. In *Proceedings of BS2015: 14th Conference of the International Building Performance Simulation Association*, 2015.
- J. Pakkanen. *Classical Greek Architectural Design: A Quantitative Approach*, volume 18 of *Papers and Monographs of the Finnish Institute at Athens*. Foundation of the Finnish Institute at Athens, 2013. doi: 10.62444/fia.1073.
- B. C. Paulson Jr. Designing to reduce construction costs. *Journal of the Construction Division (ASCE)*, 102(CO4):587–592, 1976.
- A. Prékopa. Programming under probabilistic constraint, and maximizing a probability under constraints. In G. Hellwig, P. Kall, and P. Abel, editors, *Statistical Methods for Decision Processes*, pages 78–106. Daimler-Benz, Stuttgart, Germany, 1994.
- S. A. Pyhrr. A computer simulation model to measure the risk in real estate investment. *Real Estate Appraiser*, 39:13–19, 1973.

- T. S. N. Rachmawati and S. Kim. A risk management model of apartment development projects using system dynamics. *International Journal of Construction Management*, 23(16):2742–2753, dec 2023. doi: 10.1080/15623599.2022.2045573.
- C. Rausch, M. Nahangi, C. Haas, and W. Liang. Monte Carlo simulation for tolerance analysis in prefabrication and offsite construction. *Automation in Construction*, 103:300–314, jul 2019. doi: 10.1016/j.autcon.2019.03.026.
- Rider Levett Bucknall. RLB Construction Cost Report: North America Q1 2026. Quarterly market report, Rider Levett Bucknall, March 2026. URL <https://www.rlb.com/americas/insight/rlb-construction-cost-report-north-america-q1-2026/>.
- H. H. Rong, J. Yang, M. Kang, and A. Chegut. The value of design in real estate asset pricing. *Buildings*, 10(10):178, 2020. doi: 10.3390/buildings10100178. URL <https://doi.org/10.3390/buildings10100178>.
- S. Rosen. Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy*, 82(1):34–55, 1974. doi: 10.1086/260169. URL <https://doi.org/10.1086/260169>.
- A. M. Ross and D. E. Hastings. The tradespace exploration paradigm. In *INCOSE International Symposium*, 2005.
- RSMeans. *RSMeans Building Construction Cost Data*. Gordian, 2024. Use the specific edition/year actually consulted.
- SAH Archipedia. Streamlined Moderne. Society of Architectural Historians, 2026. URL <https://sah-archipedia.org/Styles/Streamlined-Moderne>. Reference describing horizontal lines, rounded corners, nautical railings, and porthole-window motifs.
- Seattle City Council. Floor area ratio requirements. https://www.seattle.gov/documents/Departments/Council/Members/OBrien/FAR_One-Pager.pdf, 2019. Accessed 2026-06-07.
- Seattle Department of Construction and Inspections. Mandatory Housing Affordability (MHA) Program. City of Seattle, 2026. URL [https://www.seattle.gov/sdci/codes/codes-we-enforce-\(a-z\)/mandatory-housing-affordability-\(mha\)-program](https://www.seattle.gov/sdci/codes/codes-we-enforce-(a-z)/mandatory-housing-affordability-(mha)-program). Official SDCI code page for Seattle’s MHA program.
- Seattle Office of Housing. Mandatory Housing Affordability. City of Seattle, 2026. URL <https://www.seattle.gov/housing/housing-developers/mandatory-housing-affordability>. Official Office of Housing page explaining MHA performance and payment options.

- Seattle Office of Planning and Community Development. One Seattle Plan: Comprehensive Plan. City of Seattle, 2025. URL <https://www.seattle.gov/opcd/one-seattle-plan>. Official city comprehensive plan page describing the plan as Seattle's 20-year framework for growth, housing, transportation, and public investment.
- M. K. Sharma. Monte carlo simulation applications for construction project management. *International Journal of Civil Engineering and Technology*, 11(2):88–100, 2020. Journal discontinued from Scopus; consider a higher-reputation substitute.
- D. K. Sobek, A. C. Ward, and J. K. Liker. Toyota's principles of set-based concurrent engineering. *Sloan Management Review*, 40(2):67–83, 1999.
- Solemma. Custom daylight availability. <https://climatestudiodocs.com/docs/daylightCustom.html>, 2024. ClimateStudio documentation. Accessed 2026-06-07.
- W. G. Stillwell, D. A. Seaver, and W. Edwards. A comparison of weight approximation techniques in multiattribute utility decision making. *Organizational Behavior and Human Performance*, 28(1):62–77, 1981.
- Tamarkin Company. 10 Sullivan. <http://tamarkinco.com/10-sullivan>, 2026. Accessed June 10, 2026.
- S. Titman. Urban land prices under uncertainty. *American Economic Review*, 75(3):505–514, 1985. URL <https://www.jstor.org/stable/1814815>.
- A. Touran and E. P. Wiser. Monte carlo technique with correlated random variables. *Journal of Construction Engineering and Management*, 118(2):258–272, 1992. doi: 10.1061/(ASCE)0733-9364(1992)118:2(258).
- United Nations Development Programme. *Human Development Report 2010 — The Real Wealth of Nations: Pathways to Human Development*. Palgrave Macmillan, New York, 2010.
- Urban Land Institute and PwC. Emerging trends in real estate: United states and canada 2024. Technical report, Urban Land Institute, 2024.
- U.S. Department of Housing and Urban Development. Revised method for calculating the cost not attributable to dwelling use. <https://www.hud.gov/sites/dfiles/OCHCO/documents/2023-16hsgml.pdf>, 2023. Mortgagee Letter 2023-16. Accessed 2026-06-07.
- M. van Bakergem. *Exploring Health and Well-Being Considerations in Multifamily Real Estate Development*. PhD thesis, TODO: Verify institution, 2020. Doctoral dissertation. Chapter cites pp. 117–123 and 143–167.

- X.-x. Wang and J.-w. Huang. Risk analysis of construction schedule based on monte carlo simulation. In *Proceedings of the IEEE International Conference on Engineering Management*, 2009. Verify exact venue/title.
- A. C. Ward, J. K. Liker, J. J. Cristiano, and D. K. Sobek. The second toyota paradox: How delaying decisions can make better cars faster. *Sloan Management Review*, 36(3):43–61, 1995.
- A. Warszawski and R. Sacks. Practical multifactor approach to evaluating risk of investment in engineering projects. *Journal of Construction Engineering and Management*, 129(3):254–262, 2003. doi: 10.1061/(ASCE)0733-9364(2004)130:3(357).
- R. E. Weber, C. Mueller, and C. Reinhart. A hypergraph model shows the carbon reduction potential of effective space use in housing. *Nature Communications*, 15:8334, 2024. doi: 10.1038/s41467-024-52506-z. URL <https://doi.org/10.1038/s41467-024-52506-z>.
- Wheeler Kearns Architects. The construction trilemma. <https://wkarch.com/the-construction-trilemma-2/>, 2014. Practitioner commentary; grey source.
- S. K. Wong, K. W. Chau, Y. Yau, and A. K. C. Cheung. Property price gradients: The vertical dimension. *Journal of Housing and the Built Environment*, 26(1): 33–50, 2011. doi: 10.1007/s10901-010-9203-8.
- Yardi Matrix. Seattle Multifamily Market Report, February 2026. Yardi Matrix, 2026. URL <https://www.yardimatrix.com/blog/seattle-multifamily-market-report/>. Public Seattle multifamily report used as a current-market cross-check for rent and occupancy conditions.
- Y. Yoo, S. W. Hong, and J. K. Lee. Definition of the architectural style metric: An approach to quantitative analysis of design using language-image model. *Frontiers of Architectural Research*, 2025. doi: 10.1016/j.foar.2025.08.002. URL <https://doi.org/10.1016/j.foar.2025.08.002>. Accepted/in press. Update year/volume/pages if final publication data changes.

7 Appendix

7.1 Variable and Symbol Index

Symbol	Definition
A_{core}	Gross area occupied by the service core on a representative floor.
$A_{\text{core},i}$	Gross area occupied by the service core at level i , used when floor-specific core area is available.
$A_{\text{floor},i}$	Gross floor-plate area at level i .
$\bar{A}_{\text{floor},i}$	Average floor-plate area associated with design state X_i , used in the baseline-relative aspect ratio.
$\bar{A}_{\text{floor},0}$	Average floor-plate area of the baseline massing X_0 .
$A_{\text{lease},i}$	Measured leasable area at level i .
$A_{\text{net},i}$	Net leasable area at level i .
$A_{\text{sell},i}$	Sellable residential area at level i , used in condominium-tenure revenue calculations where applicable.
A_{site}	Fixed site area used to normalize programmatic yield and effective FAR.
A_{skin}	Exterior envelope surface area of the massing.
C or $\mathbb{C}$	Final cost-facing Architectural Complexity Index used in the financial model.
$C(X_i)$	Final Architectural Complexity Index associated with design state or authored iteration X_i .
$C_G(X_i)$	Pre-height geometric Complexity score for design state X_i , assembled from normalized envelope, stack, edge, and core proxies.
$C_T(X_i)$	Height-exposed cost-facing transfer score derived from $C_G(X_i)$ and the baseline-relative aspect ratio.
C_{env}	Envelope Burden; raw surface-to-volume proxy measuring exterior envelope area relative to enclosed building volume.
$\hat{C}_{\text{env}}$	Normalized Envelope Burden score.
C_{stack}	Stack Variation; raw coefficient of variation of gross floor-plate area across levels.
$\hat{C}_{\text{stack}}$	Normalized Stack Variation score.

Symbol	Definition
C_{edge}	Edge Variation; raw coefficient of variation of exterior floor-plate perimeter across levels.
$\hat{C}_{\text{edge}}$	Normalized Edge Variation score.
C_{core}	Core Offset; raw average normalized displacement between the floor-plate centroid and the service-core centroid.
$\hat{C}_{\text{core}}$	Normalized Core Offset score.
$c_j(X_i)$	Raw nonnegative proxy j extracted from design state X_i .
$\hat{c}_j(X_i)$	Normalized proxy j in the interval $[0, 1)$, produced by the reference-saturation function.
$\mathbf{c}_{\text{core}}$	Centroid of the service-core footprint.
$\mathbf{c}_{\text{floor},i}$	Centroid of floor plate i .
Δ	Grid resolution used in the stochastic design sweep.
ΔX_i	Design transition from iteration X_i to iteration X_{i+1} .
$\Delta \text{DSCR}_{i,j}$	Change in DSCR between two successive design states under paired market draw ω_j .
$\Delta \text{IRR}_{i,j}$	Change in IRR between two successive design states under paired market draw ω_j .
$\Delta \text{NOI}_{i,j}$	Change in NOI between two successive design states under paired market draw ω_j .
$\Delta \text{Robust}_{i,j}$	Change in the Boolean robustness result between two successive design states under paired market draw ω_j .
$\Delta \text{TDC}_{i,j}$	Change in total development cost between two successive design states under paired market draw ω_j .
$\Delta \text{YOC}_{i,j}$	Change in yield-on-cost between two successive design states under paired market draw ω_j .
DSCR	Debt Service Coverage Ratio.
DSCR_i	Debt Service Coverage Ratio in Monte Carlo iteration i or for design state X_i , depending on context.
$\mathcal{E}$	Financial Efficacy score used as a deterministic quality-to-complexity screen.
$\mathcal{E}_0$	Baseline Efficacy score associated with reference massing X_0 .
$\mathcal{E}_i$	Efficacy score associated with design state or authored iteration X_i .

Symbol	Definition
$\mathcal{E}_{\text{rel},i}$	Baseline-relative Efficacy score for design state X_i , defined as $\mathcal{E}_i/\mathcal{E}_0$.
F	Total number of floors in a design state.
FAR	Floor Area Ratio, equal to gross floor area divided by site area.
GFA	Gross Floor Area.
GFA_0	Baseline gross floor area used to calibrate the Efficacy constants.
GFA_i	Gross floor area associated with design state or authored iteration X_i .
H_i	Height or story-count measure associated with design state X_i .
H_0	Height or story-count measure of the baseline massing X_0 .
IRR	Internal Rate of Return.
IRR_i	Internal Rate of Return in Monte Carlo iteration i or for design state X_i , depending on context.
k_i	Half-diagonal of the floor bounding box at level i , used to normalize core offset.
κ_T	Site-topography cost coefficient translating terrain complexity into a proportional site-condition factor.
L_i	Length of exterior perimeter segment i used in the solar-orientation calculation.
Λ_K	Baseline capital-cost constant in the Efficacy index, defined as annualized baseline development cost per gross square foot.
Λ_R	Baseline rental-yield constant in the Efficacy index, defined as baseline NOI per gross square foot.
LTC	Loan-to-Cost ratio.
M_{dev}	Developer profit margin in condominium-tenure testing, defined relative to total development cost after sale proceeds and sales costs are applied.
μ	Mean or central tendency of a sampled stochastic variable.
N	Number of Monte Carlo simulation iterations.
n	Number of exterior perimeter segments used in the solar-orientation calculation.

Symbol	Definition
NOI	Net Operating Income.
NOI_0	Baseline Net Operating Income used to calibrate Λ_R .
Ω	Stochastic market-state vector containing exogenous market and financial variables sampled in the Monte Carlo simulation.
$\Omega^{(i)}$	One sampled realization of the stochastic market-state vector in Monte Carlo iteration i .
ω_j	Paired market draw used to compare design alternatives under common random numbers.
p_{delay}	Probability of a major delay event in the stochastic market vector.
$P(\Delta\text{DSCR} > 0)$	Probability that a design transition improves DSCR across paired market draws.
$P(\Delta\text{IRR} > 0)$	Probability that a design transition improves IRR across paired market draws.
$P(\Delta\text{NOI} > 0)$	Probability that a design transition improves NOI across paired market draws.
$P(\Delta\text{YOC} > 0)$	Probability that a design transition improves yield-on-cost across paired market draws.
$P(\text{IRR} \geq \tau_{\text{IRR}})$	Probability that a design state clears the investor return hurdle.
$P(\text{Robust})$	Probability that a design state clears the governing robustness test.
$P(\text{Robust} \mid C, Q)$	Probability that a design state at coordinate (C, Q) clears the governing robustness thresholds.
P_i	Exterior floor-plate perimeter at level i .
P_{10}	Tenth percentile of an output distribution.
P_{25}	Twenty-fifth percentile of an output distribution.
P_{50}	Fiftieth percentile, or median, of an output distribution.
P_{75}	Seventy-fifth percentile of an output distribution.
P_{90}	Ninetieth percentile of an output distribution.
φ	Golden-ratio exponent used in the Efficacy index, with $\varphi = (1 + \sqrt{5})/2$.
Q or $\mathbb{Q}$	Reporting Asset Quality Index and quality ordinate of the Safe Design Envelope.

Symbol	Definition
$Q(X_i)$	Reporting Asset Quality Index associated with design state or authored iteration X_i .
$Q_A(X_i)$	Area-facing quality transfer score for design state X_i , derived primarily from yield and spatial efficiency.
$Q_R(X_i)$	Rate-facing quality transfer score for design state X_i , derived from solar exposure and qualitative strategic scores.
$Q_{\text{geo}}(X_i)$	Pre-height geometric quality score for design state X_i , assembled from yield, efficiency, and solar exposure.
$Q_{\text{qual}}(X_i)$	Qualitative quality score for design state X_i , assembled from green, innovation, and material scores.
Q_{yield}	Yield; raw programmatic yield proxy measuring average net leasable area per floor relative to fixed site area.
$\hat{Q}_{\text{yield}}$	Normalized Yield score.
Q_{eff}	Spatial Efficiency; raw net-to-gross efficiency proxy measuring leasable area relative to gross floor area.
$\hat{Q}_{\text{eff}}$	Normalized Spatial Efficiency score.
Q_{sun}	Solar Exposure; raw length-weighted orientation score of the exterior perimeter.
$\hat{Q}_{\text{sun}}$	Normalized Solar Exposure score.
Q_{green}	Green Score; designer-assigned score for environmental ambition, landscape integration, high-performance systems, and related sustainability strategy.
$\hat{Q}_{\text{green}}$	Normalized Green Score.
Q_{innov}	Innovation Score; designer-assigned score for advanced construction methods, spatial differentiation, programmatic novelty, or market-positioning innovation.
$\hat{Q}_{\text{innov}}$	Normalized Innovation Score.
Q_{mat}	Material Score; designer-assigned score for material quality, tectonic ambition, facade depth, structural expression, and assembly specificity.
$\hat{Q}_{\text{mat}}$	Normalized Material Score.
q_{prem}	Market-level ceiling for the rent premium attributable to design quality in apartment-tenure calculations.

Symbol	Definition
q_{sale}	Market-level ceiling for the sale-price premium attributable to design quality in condominium-tenure calculations.
r_j	Reference value for proxy j , where $\hat{c}_j = 0.5$.
Robust_i	Boolean robustness condition for Monte Carlo iteration i . In apartment tenure, robustness requires the IRR and DSCR thresholds to clear; in condominium tenure, robustness requires the IRR and developer-margin thresholds to clear.
S_{gross}	Gross condominium sale proceeds before sales costs.
S_{net}	Net condominium sale proceeds after sales costs and development costs are applied.
s_i	Orientation score for perimeter segment i , defined as $s_i = \theta_i / \pi$.
SDE	Safe Design Envelope; the region of (C, Q) space where robustness probability exceeds the required threshold τ_P .
σ	Standard deviation or spread parameter of a sampled stochastic variable.
T	Site-topography complexity index, bounded in $[0, 1]$, used as an exogenous site-condition factor.
TDC	Total Development Cost.
TDC_0	Baseline Total Development Cost used to calibrate Λ_K .
TDC_i	Total Development Cost associated with design state or authored iteration X_i .
θ_i	Angular distance between perimeter segment i 's outward normal and true north.
τ	Investor hurdle rate used to annualize baseline capital cost in the Efficacy formulation.
τ_{DSCR}	Minimum lender DSCR threshold used in apartment-tenure robustness testing.
τ_{IRR}	Investor return hurdle.
τ_M	Minimum developer-margin threshold used in condominium-tenure robustness testing.
τ_P	Minimum robustness probability required for a design state to fall inside the Safe Design Envelope.
V	Enclosed building volume.

Symbol	Definition
v^*	Break-even rent premium required to satisfy the governing feasibility threshold.
v_{DSCR}^*	Break-even rent premium required to satisfy the DSCR threshold.
v_{IRR}^*	Break-even rent premium required to satisfy the IRR threshold.
X_0	Baseline massing used as the reference condition for relative comparison.
X_i	Design state or authored iteration i .
Y	Project outcome variable expressed as a probabilistic response to C , Q , and Ω .
YOC	Yield-on-Cost, equal to stabilized NOI divided by total development cost.
χ_i	Baseline-relative aspect ratio for design state X_i , measuring height and slenderness exposure relative to the baseline.