POLITECNICO DI TORINO

COLLEGE OF ARCHITECTURE AND DESIGN
DEPARTMENT OF ARCHITECTURE AND DESIGN (DAD)

Master's degree in
Architecture Construction City

Master Thesis

A Methodology for AI-Assisted Real-Estate Decision-Making

An auditable workflow based on SSOT, gated deliverables, and human accountability,
Case study: SPINA 3 – Corso Principe Oddone, Turin (Italy)

Supervisor:
Prof. Marta Carla Bottero

Co-supervisor:
Prof. Federico Dell'Anna

Student:
Tian Haoge

Academic Year: 2025/2026

Abstract

Early-stage real-estate development decisions bring together planning and regulatory constraints, spatial feasibility, market positioning, cost and schedule logic, and financial risk under uncertainty. In the Italian/EU context, this complexity is reinforced by demanding approval processes, fragmented local data, and strong expectations for documentation and accountability. Recent advances in artificial intelligence (AI), especially large language models (LLMs), suggest a faster way to organize feasibility work and compare alternative options. However, practical use is still limited by hallucinations, inconsistent outputs across iterations, privacy concerns, and the difficulty of producing spatially verifiable results.

This thesis investigates to what extent AI can support an end-to-end early-stage feasibility and decision-making process for a complex urban redevelopment project in Italy, and which governance mechanisms are required to make the outputs reliable, auditable, and usable in practice. Building on a review of traditional feasibility workflows and AI tool ecosystems (Chapters 2–3), the study proposes an AI-integrated methodology structured around curated evidence packaging, a Single Source of Truth (SSOT) for numeric assumptions, explicit freeze points, gate-based deliverables with acceptance criteria, and clear responsibility boundaries supported by human audit checkpoints (Chapter 4).

The methodology is applied to the SPINA 3 – Corso Principe Oddone case study in Turin (Chapter 5). The empirical AI-assisted workflow produces three comparable concept options and a versioned artefact chain spanning concept definition, parametric feasibility checks through Autodesk Forma, base-case financial modelling, sensitivity and scenario testing, and probabilistic risk analysis. A generative image model is used only as a downstream communication layer to improve figure quality, while the underlying values remain locked to the SSOT.

The findings show that single AI tool "end-to-end" support is not realistic for multi-domain redevelopment decisions. Nevertheless, a governed LLM-centered tool chain can accelerate and structure feasibility work when assumptions are frozen, evidence and numbers remain traceable, and manual audit checkpoints are treated as mandatory. Two persistent bottlenecks remain: spatial "truth", which depends on geometry and rules and therefore requires parametric validation rather than text inference; and structured-output reliability, especially formula-bearing spreadsheets and consistent cross-references, which shapes time cost and motivates session segmentation and patch-style updates. Overall, the thesis frames valuation and economic evaluation as a design-support bridge that accompanies spatial decision-making, rather than ex-post control, while positioning AI as an accelerator and integrator under governance—not a substitute for professional judgement, accountability, or stakeholder negotiation. As AI capabilities continue to evolve, this suggests that the most durable value in architectural practice will lie in evidence curation, responsible appraisal, and the integration of spatial and feasibility reasoning within transparent decision frameworks.

Keywords: project appraisal; real-estate feasibility; AI-assisted decision-making; large language models; spatial design support; workflow governance; SSOT; parametric modelling; Turin.

Table of Contents

Abstract 1

List of Figures 6

List of Tables 9

1. Introduction 11

1.1 Background: feasibility as a multi-domain decision problem 11

1.2 Motivation: Why AI is attractive, and why it is risky 11

1.3 Research gaps and positioning of this thesis 12

1.4 Research questions 12

1.5 Objectives and expected contributions 13

1.6 Methodology overview and scope 13

1.7 Structure of the thesis 14

2. Traditional Real-Estate Evaluation and Early-Stage Design Processes 15

2.1 Market analysis and demand/supply assessment 15

2.2 Economic and financial evaluation 17

2.3 Planning constraints, regulatory framework, and preliminary design 20

2.4 Limitations of the traditional workflow 22

3. Literature Review: AI and Digital Tools in Real-Estate Decision-Making 25

3.1 Overview of AI developments relevant to planning, design and feasibility 25

3.2 Applications of AI in architecture, planning and real estate 26

3.2.1 Market and demand analytics 27

3.2.2 Valuation and risk assessment 27

3.2.3 Design and planning decision support 28

3.2.4 Participation, communication, and document work 28

3.3 Digital and data-driven tools in feasibility and early-stage design 28

3.3.1 BIM, CDEs, and interoperability 28

3.3.2 GIS, urban datasets, remote sensing, and digital twins 29

3.3.3 Parametric design and early-stage simulation 29

3.3.4 Financial modelling infrastructure and model risk 29

3.4 SWOT analysis of available AI and digital tools 30

3.4.1 SWOT Analysis of Vertical AI and Digital Tools 31

3.4.2 SWOT Analysis of Large Language Models 40

3.5 Research gaps 43

4. Integration of AI into the Real-Estate Decision-Making 45

4.1 Basic Process 45

4.2 Data Collection and Upload 46

4.2.1 Data Collection Checklist 47

4.2.2 Principles for Data Organization and Upload 49

4.3 Testing and Selection of Large Language Models 49

4.3.1 Large Language Model Testing Methodology 50

4.3.2 Testing Procedures for Major Large Language Models 50

4.3.3 Scoring Sheet and Scoring Results Discussion 53

4.4 Final Methodology for Integration of AI into the Real-Estate Decision-Making, Based on Test Scores and SWOT Analysis 55

5. Empirical Application to a Real Case Study 57

5.1 Case Study Introduction 57

5.2 Workflow Logic and Principle 60

5.3 Workflow of Methodology for AI Integration Across the Entire Real Estate Decision-Making Process 61

5.3.1 Concept Report 61

5.3.2 Modeling and Technical Feasibility Study 62

5.3.3 Financial Modeling, Sensitivity Analysis, and Risk Assessment 69

5.3.4 Integrated Comparative Decision-Making 86

5.3.5 Supplement with Visual Charts 87

5.3.6 Manual Adjustments and Final Delivery 93

5.4 How AI Decision Reports Support Spatial Iteration: Translating Feasibility Logic into a Phased Plan, Forma Massing, and Communication Outputs 94

5.4.1 Step A - Extract decision-relevant information from the AI report (not geometry) 94

5.4.2 Step B - Lock the anchor condition and the "base concept" (core assumptions) 94

5.4.3 Step C - Translate the report into a "new planning scheme" that is spatially testable . 94

5.4.4 Step D - Convert the scheme into a phased delivery plan (Phase 0–3 with explicit scope and triggers) (Revised as requested) 95

5.4.5 Step E - Model in Autodesk Forma: keep massing aligned with the phased plan 97

5.4.6 Step F - Communication outputs: stylised master plan and context compositing (AI as visual assistant, not SSOT) 99

5.5 End-to-end Workflow Summary and Lessons Learned 101

6. Discussion 106

6.1 Critical reflection on the integration of AI into planning and development feasibility 106

6.2 Strengths and weaknesses observed 109

6.3 Methodological implications 111

6.4 Forward-looking implications for professional practice and the real-estate sector (2025-2026 horizon) 113

7. Conclusions 115

7.1 Answers to research questions 115

7.2 Key findings 116

7.3 Limitations 117

7.4 Looking ahead: AI acceleration and the role of architectural practitioners 117

7.5 Recommendations for practice and future research 118

Attachments 119

References 169

References (Chapters 2–3) 169

References (Chapter 4) 171

References (Chapter 6) 172

List of Figures

Figure 3.1. Timeline / typology of AI developments and typical built environment use cases.

Figure 4.1. AI-integrated feasibility workflow for early-stage real-estate decision-making (methodological diagram).

Figure 4.2. Overall model score comparison (weighted average across two decision-report runs).

Figure 4.3. Dimension-profile comparison across models (average of two runs per model).

Figure 5.1 Indicative satellite view and boundary of the "Spina 3 – Oddone" site along Corso Principe Oddone (Turin). Source: Italian Trade Agency (ICE), 2016.

Figure 5.2 Bird's-eye view context of Corso Principe Oddone and Piazza Baldissera, illustrating the infrastructural character and urban edges. Source: Città di Torino, 2023.

Figure 5.3 Two-chat workflow for the empirical case study, showing an input-side prompt/data session and an output-side generation session, with a human audit checkpoint that validates each versioned artefact before proceeding. Source: Author's diagram, 2025.

Figure 5.4 Autodesk Forma massing models for Proposals 1, shown at consistent scale and viewpoint for comparative review.

Please refer to the attachment for the report without human edition.Figure 5.5 Autodesk Forma massing models for Proposals 2, shown at consistent scale and viewpoint for comparative review.

Figure 5.6 Autodesk Forma massing models for Proposals 3, shown at consistent scale and viewpoint for comparative review.

Figure 5.7 Example Autodesk Forma environmental-analysis dashboard outputs for Proposal 3, illustrating concept-stage comparative checks (Sun hours). Source: Author's Autodesk Forma outputs, 2025.

Figure 5.8 Example Autodesk Forma environmental-analysis dashboard outputs for Proposal 3, illustrating concept-stage comparative checks (Daylight potential). Source: Author's Autodesk Forma outputs, 2025.

Figure 5.9 Example Autodesk Forma environmental-analysis dashboard outputs for Proposal 3, illustrating concept-stage comparative checks (Embodied carbon). Source: Author's Autodesk Forma outputs, 2025.

Figure 5.10 Example Autodesk Forma environmental-analysis dashboard outputs for Proposal 3, illustrating concept-stage comparative checks (Microclimate). Source: Author's Autodesk Forma outputs, 2025.

Figure 5.11 Example Autodesk Forma environmental-analysis dashboard outputs for Proposal 3, illustrating concept-stage comparative checks (Solar energy). Source: Author's Autodesk Forma outputs, 2025.

Figure 5.12 Example Autodesk Forma environmental-analysis dashboard outputs for Proposal 3, illustrating concept-stage comparative checks (Wind). Source: Author's Autodesk Forma outputs, 2025.

Figure 5.13 Tornado sensitivity chart – Concept 1 (Gate 4 output). Source: SPINA3_Sensitivity_Tornado_v1.2.2_20251227_2pagesPerConcept_CleanFilled.pdf, 2025 (pp. 1–6).

Figure 5.14 Tornado sensitivity chart – Concept 2 (Gate 4 output). Source: SPINA3_Sensitivity_Tornado_v1.2.2_20251227_2pagesPerConcept_CleanFilled.pdf, 2025 (pp. 1–6).

Figure 5.15 Tornado sensitivity chart – Concept 3 (Gate 4 output). Source: SPINA3_Sensitivity_Tornado_v1.2.2_20251227_2pagesPerConcept_CleanFilled.pdf, 2025 (pp. 1–6).

Figure 5.16 3×3 scenario matrix visualisation (Market × Delivery shocks): Project NPV (EUR m) for Concepts 1–3 (Gate 4 output; derived from SPINA3_Scenario_Matrix_v1.0.1_20251227_Viz_v1.xlsx, sheets Matrix_C1–Matrix_C3). Source: SPINA3_Scenario_Matrix_v1.0.1_20251227_Viz_v1.xlsx, 2025 (sheets Matrix_C1–Matrix_C3).

Figure 5.17 Risk heat map (Probability × Impact; counts by score pair). Source: SPINA3_RiskRegister_v1.0.4_20251227.xlsx, RiskRegister sheet (Probability_Score_1to5, Impact_Score_1to5; derived count matrix).

Figure 5.18 Monte Carlo distributions (NPV and IRR histograms) for C1–C3. Source: SPINA3_MonteCarlo_Results_C1-C3_v1.0.2_20251227.xlsx, Sim_C1/Sim_C2/Sim_C3 (NPV, IRR columns).

Figure 5.19 Multi-concept Gantt comparison. Source: Author's Nano Banana Pro outputs, 2025.

Figure 5.20 Cash-flow curves (C1–C3). Source: Author's Nano Banana Pro outputs, 2025.

Figure 5.21 Scenario matrix visualization (3×3). Source: Author's Nano Banana Pro outputs, 2025.

Figure 5.22 Tornado sensitivity charts. Source: Author's Nano Banana Pro outputs, 2025.

Figure 5.23 Monte Carlo distribution and CDF plots. Source: Author's Nano Banana Pro outputs, 2025.

Figure 5.24 Phase 0 baseline (empty-site control map). Source: Author, exported from Autodesk Forma.
Figure 5.25 Phase 1 plan (residential + commercial + dormitory + institutional anchor + initial public space). Source: Author, exported from Autodesk Forma.
Figure 5.26 Phase 2 plan (residential + commercial + co-working/office + expanded public space). Source: Author, exported from Autodesk Forma.
Figure 5.27 Phase 3 plan (completion: low-density housing + flexible infill + public-space continuity). Source: Author, exported from Autodesk Forma.
Figure 5.28 Autodesk Forma bird's-eye output (massing + context). Source: Author, Autodesk Forma export.
Figure 5.29 Stylized master plan (competition-oriented graphic; context de-emphasized). Source: Author, generated with Nano Banana Pro from Forma export.

Figure 5.30 Context composite (proposal inserted into realistic bird’s-eye scene). Source: Author, generated with Nano Banana Pro from base aerial + Forma massing.

Figure 5.31 End-to-end empirical workflow (two-chat structure, audit checkpoints, and versioned artefacts). Source: Author’s diagram, 2026.

List of Tables

Table 2.1. Traditional Market Analysis: Inputs → Methods → Outputs → Typical Failure Modes

Table 2.2. Financial Evaluation Toolkit: Method → When used → Key outputs → Known limitations

Table 2.3. Administrative/Planning Feasibility Checklist (Traditional) + Typical bottlenecks + consequences for schedule/cost

Table 3.1. Digital tool stack in early-stage feasibility: purpose → typical outputs → common pitfalls → governance needs.

Table 3.2. Evaluation criteria used in the SWOT: criteria → why it matters → supporting literature/standard.

Table 3.3. Practical usability filter scoring table (Author's evaluation, Dec 2025).

Table 3.4. SWOT table (Stage 1) — Market / data and benchmarking tools (Author's evaluation, Dec 2025).

Table 3.5. SWOT table (Stage 2) — Planning / early design tools (Author's evaluation, Dec 2025).

Table 3.6. SWOT table (Stage 3) — Financial feasibility and risk tools (Author's evaluation, Dec 2025).

Table 3.7. SWOT table (Stage 4) — Integrated / combined approaches (Author's evaluation, Dec 2025).

Table 3.8. SWOT of large language models for urban redevelopment decision-making (Author's evaluation, Dec 2025).

Table 4.1. Data collection checklist for the Turin case study (SPINA 3 – Corso Principe Oddone).

Table 4.2. Practical constraints observed during LLM stress testing (platform features and limits).

Table 4.3. Rubric dimensions and weights used for scoring the decision reports.

Table 4.4. Overall ranking and model scores from the stress-test rubric (average across two reports per model).

Table 5.1 SSOT summary of Concept 1–3 (program intent, key indicators, and defining features). Source: Feasibility Study – Constraints, Market & Concepts (2025).

Table 5.2 Autodesk Forma area metrics for Proposals 1–3 (exported from Area metrics; values in m2 unless noted). Source: Author's Autodesk Forma export, 2025 (Area metrics).

Table 5.3 Summary of the most decision-critical locked constraints (non-negotiable). Source: SPINA3_Inputs_HardConstraints_All_v1.0_20251224.xlsx, 2025.

Table 5.4 Concept 1–3 program mix and phasing summary (as extracted from the program-and-phasing register). Source: SPINA3_Inputs_ConceptProgramPhasing_C1-C3_v1.1.2_20251226.xlsx, 2025.

Table 5.5 Master Inputs snapshot (selected key assumptions used downstream). Source: SPINA3_Inputs_MasterInputs_All_v1.4.2.1_C2GFAFix_TraceHotfix_20251227.xlsx, 2025.

Table 5.6 Data gaps and adopted ranges/distributions (model-ready selection). Source: SPINA3_Inputs_DataGapsRanges_All_v1.3_20251227.xlsx, 2025.

Table 5.7 Baseline cost and schedule summary (WBS totals, timing, and applied cash-flow curve profiles). Source: SPINA3_CostPlan_WBS_C1-C3_v1.2.1_20251226.xlsx, 2025; SPINA3_ScheduleBaseline_Gantt_C1-C3_v1.0.4.5_20251227_UltraTraceRefPatch.xlsx, 2025.

Table 5.8 Base-case KPI outputs for Concept 1–3 (Gate 3 financial model; pre-tax, project-level). Source: SPINA3_FinancialModel_C1-C3_v1.0_20251227.xlsx, 2025.

Table 5.9 Top sensitivity drivers (Tornado; ranked by visual ordering of the Gate 4 output). Source: SPINA3_Sensitivity_Tornado_v1.2.2_20251227_2pagesPerConcept_CleanFilled.pdf, 2025.

Table 5.10 Scenario matrix results (diagonal cases: Upside_M/Upside_D, Base_M/Base_D, Downside_M/Downside_D). Source: SPINA3_Scenario_Matrix_v1.0.1_20251227_Viz_v1.xlsx, 2025.

Table 5.11 Top 10 risks by severity score (ID, category, probability/impact, mitigation/response). Source: SPINA3_RiskRegister_v1.0.4_20251227.xlsx, RiskRegister sheet (rows 2–77; columns A–N; scoring fields in columns I, K, L).

Table 5.12 Monte Carlo summary statistics by concept (P10/P50/P90 for NPV and IRR; threshold probabilities; peak funding gap). Source: SPINA3_MonteCarlo_Results_C1-C3_v1.0.2_20251227.xlsx, Summary sheet (rows 2–4; columns A–AA).

Table 5.13 Decision summary matrix (Concept 1–3): feasibility stance, strengths, binding risks, conditions/mitigations, and decision implication. Source: Integrated Investment Decision Report – SPINA 3 Redevelopment (Turin, Italy) – Revised_v1.1.docx, Sections 1, 7–8 (synthesized).

Table 5.14 Visual refinement log (chart type, refinement actions, and traceable data source). Source: Author's post-processing gate (2025); underlying data sources listed per row.

Table 5.15 End-to-end artefact map. Source: Author's synthesis based on versioned workflow artefacts, 2025–2026.

1. Introduction

This chapter introduces the research problem, motivates the thesis, and defines the research questions, objectives, and scope. It positions the work at the intersection of architecture, concept-stage feasibility, and the practical governance of AI-assisted decision support for urban redevelopment.

Rather than framing AI as a standalone generator, the thesis treats AI as a set of tools embedded in a controlled workflow: evidence is curated, constraints and assumptions are stabilized in a single source of truth (SSOT), outputs are produced through versioned iterations, and each transition is validated through explicit human audit checkpoints. As demonstrated in Section 5.4, the audited decision report also functions as a translation layer from feasibility logic to spatial actions, enabling measurable massing tests and early environmental checks using Autodesk Forma as a geometry-and-quantity SSOT. The objective is not to claim final design resolution or regulatory approval, but to establish a traceable bridge from constraints and spatial logic to comparable assumptions that can be carried forward into financial modelling and decision-making.

1.1 Background: feasibility as a multi-domain decision problem

Concept-stage real-estate feasibility sits at the intersection of planning constraints, spatial feasibility, market positioning, cost and schedule logic, and financial risk. In early stages, decision-makers must determine whether a project should proceed, be re-scoped, or be paused, often with incomplete information and under time pressure. Small changes in program mix, phasing, or buildable capacity can shift the financial outcome and risk profile, while financial constraints can in turn reshape spatial priorities.

In the Italian and wider EU context, this coupling is intensified by complex approval pathways, layered regulations, and fragmented micro-market evidence. For urban redevelopment projects, feasibility work must therefore balance rigor with iteration: it should remain responsive to new evidence while preserving a clear record of assumptions, sources, and decision steps.

Within an architectural redevelopment setting, feasibility is not separate from design. Even at a conceptual level, the credibility of a feasibility narrative depends on whether spatial intentions can be tested against basic constraints (access, massing logic, phasing, and environmental exposure) and communicated in a way that stakeholders can scrutinize.

1.2 Motivation: Why AI is attractive, and why it is risky

The acceleration of AI capabilities, especially large language models (LLMs), has created clear incentives for feasibility and early-stage design work. In 2025, AI tools are shifting from experimental, breakthrough demonstrations toward broader professional adoption, largely because they can summarize complex documents, extract and structure information, draft comparative narratives, and support rapid scenario reasoning.

At the same time, feasibility is a high-stakes domain: errors in constraints, quantities, or assumptions can propagate into design decisions and financial commitments. LLM outputs can be convincing while still being wrong, and iterative use can introduce drift across versions if assumptions are not explicitly stabilized. These risks raise a practical question for built-environment practice: how can AI be used to increase productivity without undermining accountability?

This thesis responds by focusing on governance rather than novelty. It defines a workflow where AI outputs are treated as draft artefacts that must remain traceable to inputs, constrained by an SSOT, and verified through human audit checkpoints. In Section 5.4, this approach is extended into spatial iteration: the decision report translates feasibility reasoning into testable design actions, while the phased plan and Autodesk Forma model provide an authoritative baseline for geometry and quantities. Generative visual outputs are confined to communication artefacts and are checked against that baseline to avoid mistaking visual plausibility for verified geometry.

1.3 Research gaps and positioning of this thesis

The literature on AI in the built environment highlights promising applications across planning analysis, design assistance, and real-estate evaluation. However, several gaps remain between what is technically possible and what can be used responsibly in concept-stage redevelopment feasibility, particularly in an Italy/EU setting.

First, there is an integration gap: many studies and commercial tools address isolated tasks (for example, document summarization, market insights, or visual generation), whereas feasibility decisions require an end-to-end logic that connects constraints, spatial reasoning, and financial implications. Second, there is an auditability gap: feasibility workflows depend on traceability, yet AI use often lacks stable assumptions, version control, and reproducible artefact chains. Third, there is a spatial-verification gap: textual feasibility reasoning must ultimately connect to testable geometry and quantities, but this translation is often left implicit or treated as a purely design-driven step.

Finally, production constraints limit practical adoption: the quality of output depends on input preparation, data availability, and tool interoperability, while costs, regional data coverage, and platform limitations can restrict what is feasible in real projects or student research settings.

Accordingly, the contribution of this thesis is not the creation of a new AI system, but the development and validation of a minimum viable governance stack for AI-assisted feasibility. This stack combines curated inputs, a single versioned SSOT, gated deliverables, and explicit human accountability, and it is tested empirically through a real redevelopment case study in Turin.

1.4 Research questions

The thesis addresses three research questions:

RQ1: To what extent can AI support an end-to-end feasibility and early-stage development decision process for a complex urban redevelopment project in Italy?

RQ2: What governance mechanisms are required to mitigate hallucination, scope drift, and data inconsistency, while preserving auditability and accountability?

RQ3: What are the observed strengths, limitations, and role boundaries of an LLM-centered workflow in the built-environment feasibility domain?

1.5 Objectives and expected contributions

To answer these questions, the thesis pursues five objectives:

1. Define the decision tasks and common failure modes of traditional concept-stage feasibility in an Italy/EU context, with emphasis on evidence quality, iteration costs, and sensitivity to key assumptions.
2. Review AI and digital tools relevant to feasibility and early-stage redevelopment, and evaluate their practical usability constraints (data coverage, cost, integration effort, and accountability requirements).
3. Develop an AI-integration methodology based on evidence packaging, a single source of truth (SSOT), gated deliverables, and human audit checkpoints that stabilize assumptions across iterations.
4. Apply the methodology to the Turin SPINA 3 – Corso Principe Oddone case to produce a traceable feasibility-to-decision artefact chain, including an audited decision report and a disciplined translation to measurable spatial tests and early environmental checks (Section 5.4), aligned with subsequent financial and risk modelling.
5. Synthesize lessons learned from the empirical application, clarifying where AI adds value, where it remains unreliable, and what governance conditions are necessary for responsible use in practice.

The expected contributions are methodological and operational: a replicable governance approach for AI-assisted feasibility, an auditable artefact chain that supports comparability across concept options, and a practical account of limitations and controls that can inform both academic work and professional feasibility practice.

1.6 Methodology overview and scope

The research design combines (i) a review of traditional feasibility workflows and their limitations, (ii) a literature review and practical testing of AI and digital tools, and (iii) an empirical workflow application to a real redevelopment case. Based on these steps, the thesis develops an AI-integration methodology that separates an input-side stage (curated dataset and prompt logic) from an output-side stage (generation of downstream artefacts), with a versioned SSOT and gate-based audit checkpoints that manage iteration and reduce drift.

The scope is deliberately limited to concept-stage feasibility and early design validation. The thesis does not claim full regulatory compliance, detailed design resolution, or investment-grade valuation.

Instead, it focuses on establishing a transparent decision record: how constraints and evidence are consolidated, how assumptions are stabilized, how outputs are verified, and how feasibility reasoning can be translated into spatial tests that remain anchored to a geometry-and-quantity baseline.

1.7 Structure of the thesis

The thesis is organized in a sequence that moves from problem definition to method development and empirical validation. After this introduction, the next chapter reviews traditional feasibility and early-stage design processes and summarizes their key limitations. The literature review chapter then surveys AI and digital tools relevant to the built environment and identifies practical gaps for responsible adoption. The methodology chapter formalizes the proposed governance approach, including SSOT discipline, gate-based deliverables, and audit checkpoints. The empirical chapter applies the method to the Turin SPINA 3 case, producing a traceable artefact chain from curated evidence to decision outputs and spatial testing. The discussion chapter interprets the results, examines limitations and failure patterns, and situates the findings within the broader trajectory of AI adoption in the real-estate sector. The final chapter answers the research questions, summarizes contributions, and outlines implications for practice and further research.

2. Traditional Real-Estate Evaluation and Early-Stage Design Processes

This chapter explains how real-estate projects are commonly assessed in early stages using market studies, feasibility and valuation models, and planning/administrative checks. These steps form the baseline workflow used by developers, consultants, lenders, and public authorities (Manganelli, 2015, pp. 89–92; Miles et al., 2015).

Traditional workflows persist because they produce familiar outputs—such as market comparables, cash-flow projections, and permit checklists—that support accountability and decision-making. At the same time, early-stage work often starts with incomplete local data, uncertain approval timelines, and changing design inputs, which can weaken the reliability of results (Manganelli, 2015, p. 36; Miles et al., 2015).

Key terms are used consistently throughout. Feasibility means whether a project is viable under constraints (market, cost, time, and regulation). Valuation means an estimate of value for a defined purpose and scope (Manganelli, 2015, pp. 89–91; International Valuation Standards Council [IVSC], 2024).

Risk refers to uncertainty that can be measured or reasonably quantified. Uncertainty refers to conditions that cannot be measured well, especially at early stages. This distinction matters because many models produce precise numbers even when the underlying uncertainty is high (Manganelli, 2015, p. 137; Royal Institution of Chartered Surveyors [RICS], 2024).

The following sections describe this baseline workflow and highlight where it tends to break down in practice (Manganelli, 2015, p. 36; IVSC, 2024).

2.1 Market analysis and demand/supply assessment

Market analysis starts from the fact that real-estate markets are not "perfect" markets. Properties are heterogeneous, location-bound, and traded with information gaps and uneven transparency, so analysts cannot assume that prices fully reflect all relevant information (Manganelli, 2015, pp. 8–9).

A practical first step is to define the relevant submarket. A submarket is the area and product segment where properties compete as close substitutes. This definition is not neutral: the boundary choice shapes which comparables are used and what "demand" means for the project (Manganelli, 2015, pp. 9–10).

The analyst should state the purpose of the market study before collecting data. A screening study may rely on broad indicators, while a feasibility- or valuation-oriented study needs more granular and better-documented evidence (Manganelli, 2015, p. 38).

In practice, market evidence usually combines secondary sources (statistics, registries, published reports, listing data) and primary sources (site observation, interviews, user surveys). Secondary sources can be delayed or poorly matched to the micro-location. Interviews can be biased because stakeholders may have incentives to overstate or understate market strength. Direct observation can reduce some biases, but it does not replace robust data (Manganelli, 2015, pp. 35–36).

A common problem in the early stages is that a large volume of data does not necessarily equate to high data quality. For example, while publicly available listing prices are easy to collect, they differ from actual transaction prices. Treating listing prices as actual transaction prices may overestimate rents and prices, leading to biased assessment results (Manganelli, 2015, p. 36).

Because thorough market research is time-consuming and costly, most workflows must be conducted in stages: initial screening → targeted research → assessment of option potential → detailed research. While this helps control workload, it also introduces other risks. For instance, assumptions made in the early stages may become inapplicable after design changes but remain in the financial model, leading to errors in the final estimate (Manganelli, 2015, pp. 37–38). Demand and supply assessment depends on the competitive framework. Analysts should distinguish viable competitors from obsolete stock and should consider future supply in the pipeline. If the pipeline is ignored, vacancy and absorption forecasts can become systematically optimistic (Manganelli, 2015, pp. 50–51).

Timing is part of market logic, not only a scheduling detail. Administrative procedures and development lead times affect when supply comes to market and how competing projects respond. If time-to-market is underestimated, the analysis may assume demand conditions that no longer apply when the project actually delivers (Manganelli, 2015, p. 75; Miles et al., 2015).

In some contexts, fine-grained geo-referenced transaction and rental evidence is difficult to obtain. Manganelli notes this problem explicitly for Italy, but the issue is broader: when micro-market data are scarce, analysts must use proxies, clearly label them, and state the limits of inference (Manganelli, 2015, p. 36).

Table 2.1. Traditional Market Analysis: Inputs → Methods → Outputs → Typical Failure Modes

Inputs	Methods	Outputs	Typical failure modes	Practical mitigation
Macro indicators (GDP, demographics, jobs)	Trend review; context comparison	Market context	Macro averages hide local differences; time lag	State limits; avoid micro conclusions
Local stock + submarket boundary	Segmentation by substitutability and location	Defined competitive arena	Boundary chosen for convenience; weak comparables	Document criteria; test alternatives
Prices/rents (transactions or quotations)	Comparable analysis; adjustments	Indicative price/rent levels	Quotations treated as deals; selection bias	Triangulate; label proxies; disclose gaps
Demand proxies (households, firms, user input)	User profiling; preference inference	Demand estimates by segment	Interview/survey bias; small samples	Combine with observation; report uncertainty
Pipeline + competing projects	Competitor set; scenario framing	Absorption/vacancy outlook	Future entrants ignored; obsolete stock included	Include pipeline scenarios; define obsolescence
Time-to-market and approvals timing	Lifecycle alignment; timing risk mapping	Timing-adjusted market view	Assuming quick delivery; ignores delays	Link timing to demand and revenues

Note. Synthesized from Manganelli (2015, pp. 35–36, 37–38, 50–51, 75) and the development-process emphasis on timing and coordination in Miles et al. (2015).

Mini summary

Market analysis is most useful when it defines the submarket clearly, matches methods to the study purpose, and cross-checks evidence from multiple sources. Its main weakness is that key input can be biased or delayed, especially at the micro-location level. If early proxies remain unchallenged, later decisions can rely on assumptions that no longer match the evolving design (Manganelli, 2015, pp. 35–36, 37–38).

Common pain points in practice

- Local transactions and rental data can be scarce or delayed, so analysts rely on proxies with uneven reliability (Manganelli, 2015, p.36).
- Listing quotations are often used as if they were transaction prices, which can push assumptions upward (Manganelli, 2015, p. 36).
- Early assumptions can "stick" even when the design changes, creating inconsistent demand and absorption logic (Manganelli, 2015, pp. 37–38).
- Pipeline supply is hard to observe and is often underestimated, leading to optimistic vacancy/absorption forecasts (Manganelli, 2015, pp. 50–51).
- Long approvals can make early market evidence stale by the time the project reaches the market (Manganelli, 2015, p. 75; Miles et al., 2015).

2.2 Economic and financial evaluation

Financial evaluation translates market and design assumptions into a cash-flow view of the project. In early stages, teams usually start with pre-feasibility screening and then refine assumptions as design, costs, and approvals become clearer (Manganelli, 2015, pp. 89–91).

Discounted cash flow (DCF) is the most common framework for linking cash flows over time to value. Net present value (NPV) measures the value created after discounting future cash flows. This approach is methodologically sound, but its reliability depends on the quality of timing, revenue, cost, and financing inputs (Manganelli, 2015, p. 97).

The time dimension of a project has a decisive impact on valuation. Delays in project progress not only mean a postponement of revenue recognition, but also lead to a further mismatch between upfront fixed expenditures and later returns. Even if nominal total profit remains unchanged, the delayed return to positive cash flow, constrained by the discounting effect, will significantly reduce the project's net present value. In other words, even with fully locked-in rental pricing and cost budgets, simply adjusting the phasing of the construction or sales cycle is enough to cause significant fluctuations in the final valuation (Manganelli, 2015, p. 97).

The Internal Rate of Return (IRR) has become a core, universally accepted metric in the industry because it condenses complex investment returns into a percentage value that is easy to benchmark horizontally. However, the effectiveness of this metric has clear boundaries, especially when comparing projects of vastly different sizes and with mismatched timelines. Using IRR as the sole evaluation dimension can easily distort investment decisions and should be considered in conjunction with the discount rate (Manganelli, 2015, pp. 124–128). The discount rate is a core element in translating risk into valuation. Theoretically, the discount rate should be rigorously set based on the risk-free rate plus a risk premium. However, in practice, to accommodate the expected valuation range, project teams sometimes pre-determine the discount rate, turning it into a result-oriented adjustment tool. This practice essentially sacrifices the transparency of the estimation process and makes the true level of risk implied behind the numbers difficult to discern (Manganelli, 2015, pp. 133–134; IVSC, 2024).

The terminal value, as the exit value at the projected end, is usually derived from the capitalization rate or rate of return and often carries a dominant weight in the final valuation result. This weighting means that the sensitivity of the final valuation lies more in the pre-determined exit price than in the actual cash flows generated during the development phase. It must be pointed out that even if a spreadsheet-based model appears logically consistent and computationally efficient, its calculated final value will lose its reference value once the exit rate of return, a core variable, deviates from empirical market support (Manganelli, 2015, p. 104; RICS, 2024).

While sensitivity analysis helps identify which key assumptions have the greatest impact on the results, its limitation lies in its inability to provide the probability distribution of the outcome. Furthermore, conventional univariate sensitivity analysis often overlooks the "interconnectedness" of risks, that is, risks are often interdependent; for example, project delays usually directly lead to cost increases. Therefore, simply equating sensitivity analysis results with "risk quantification" is highly likely to underestimate the impact of compound risks on the project (Manganelli, 2015, p. 154).

Traditional single-cash flow models are often limited by a static perspective and struggle to encompass dynamic risks. Introducing scenario analysis and decision trees can fill this gap by visualizing event sequences and interconnected relationships. However, it is undeniable that the effectiveness of such methods in the early stages of a project remains highly subjective. Whether it's the construction of the scenario framework or the allocation of probability weights, it is inherently difficult to escape the reliance on the decision-maker's experience. Therefore, rather than demanding absolute precision in the output, it is more pragmatic to define this type of analysis as a structured deductive tool to aid thinking (Manganelli, 2015, pp. 154–155).

At this stage, clearly defining the boundary between "risk" and "uncertainty" is particularly necessary. Once deep uncertainty is involved, such as uncontrollable approval cycles or drastic fluctuations in future market demand, the model's output will inevitably be influenced by its pre-existing assumptions. Therefore, clearly defining and disclosing the scope of the assessment and data boundaries is crucial. Its purpose is to warn users not to be misled by seemingly precise but actually fragile figures (Manganelli, 2015, p. 137; IVSC, 2024; RICS, 2024).

In addition, developers often use heuristics such as the "hurdle rate" to speed up the screening process and clarify the bottom line for investment to stakeholders. While this approach is quite efficient in practice, it can also easily obscure the deeper driving factors that affect the success or failure of a project—especially when a simple "meeting the criteria is enough" mentality replaces transparent discussions of risk mechanisms, its hidden dangers are particularly prominent (Moorhead, Armitage, & Skitmore, 2024).Table 2.2 summarizes the main financial tools used in early-stage evaluation. It explains when each method is used and the most common limitations that affect decision quality (Manganelli, 2015, pp. 89–91, 104, 124–128, 154).

Table 2.2. Financial Evaluation Toolkit: Method → When used → Key outputs → Known limitations

Method	When used	Key outputs	Known limitations
Pre-feasibility screening	Early option filtering	Go/no-go; rough margin	High uncertainty; proxy inputs can persist (Manganelli, 2015, pp. 89–91)
DCF (NPV)	Core feasibility/valuation with timeline	NPV; value over time	Input quality dominates; timing sensitive (Manganelli, 2015, p. 97)
IRR	Comparison across projects	IRR; return profile	Can mis-rank; multiple IRRs possible (Manganelli, 2015, pp. 124–128)
Discount-rate choice	Risk-adjusted value estimate	Discount rate; present values	May be adjusted to “fit”; weak transparency (Manganelli, 2015, pp. 133–134; IVSC, 2024)
Terminal value / exit yield	End-of-period value estimate	Exit price; yield/cap rate	Often dominates results; evidence can be thin (Manganelli, 2015, p. 104)
Sensitivity analysis	Driver testing	Key driver ranking	No probabilities; weak for correlated shocks (Manganelli, 2015, p. 154)
Scenario analysis	Plausible cases	Scenario NPVs/IRRs	Scenario design is subjective (Manganelli, 2015, p. 154)
Decision trees	Sequential decisions	Conditional NPVs	Branch weights often judgmental (Manganelli, 2015, pp. 154–155)

Note. Synthesized from Manganelli (2015, pp. 89–91, 97, 104, 124–128, 133–134, 154–155) and governance expectations in IVSC (2024) and RICS (2024).

Mini summary

Financial evaluation is effective when it links a clear timeline to consistent revenue, cost, and financing assumptions. The main vulnerability is that small changes in timing, discount rate, or exit yield can flip results. Sensitivity and scenarios help interpretation, but they are not a substitute for strong evidence and transparent disclosure (Manganelli, 2015, pp. 104, 133–134, 154; RICS, 2024).

Common pain points in practice

- Discount rates can become negotiable parameters, which reduces transparency about the true risk level (Manganelli, 2015, pp. 133–134; IVSC, 2024).
- Terminal value can dominate the result even when the exit yield is weakly supported by market evidence (Manganelli, 2015, p. 104).
- Delays often occur together with cost escalation, but one-variable sensitivities do not capture this correlation (Manganelli, 2015, p. 154).
- Spreadsheets can hide model risk through overrides, inconsistent scenarios, or weak version control (IVSC, 2024; RICS, 2024).
- IRR can mis-rank options and can distract attention from timing and scale effects (Manganelli, 2015, pp. 124–128).

2.3 Planning constraints, regulatory framework, and preliminary design

Planning and administrative reviews define the scope and timelines of construction, determining project feasibility. While procedures may vary by country, most systems include permitting and compliance reviews, which directly impact the buildable size and area, documentation workload, and project schedule. This article uses Italy as an example to illustrate permitting and compliance reviews (Miles et al., 2015; D.P.R. 380/2001, art. 12).

In Italy, major projects such as new construction and urban redevelopment require a building permit (permesso di costruire). This means feasibility studies should identify which design options warrant formal approval as early as possible, and determine the supporting documentation required for the application (D.P.R. 380/2001, art. 10). Permit issuance depends on conformity with planning instruments and building regulations. In practice, it may also depend on infrastructure and urbanization conditions, which can introduce additional steps and coordination requirements. If these obligations are discovered late, they can produce schedule shocks and unbudgeted costs (D.P.R. 380/2001, art. 12).

Urban standards also constrain preliminary design. For example, mandatory minimum distances between buildings affect massing, daylight access, and site capacity. If early massing is not checked against such rules, teams may invest in concepts that later become non-compliant and require redesign (D.M. 1444/1968, art. 9).

Approvals are rarely linear. Documentation requests, stakeholder negotiation, and revisions can create iterative cycles. This matters because time is a cost: delays reduce discounted value and can expose the project to different market conditions (Miles et al., 2015; Manganelli, 2015, p. 75).

In the early stages of a project, the accuracy of cost data is often limited by insufficient design depth. At this point, classifying and grading estimates becomes particularly necessary. By establishing a correlation between design depth and expected accuracy, it provides a theoretical benchmark for defining the effective boundaries of estimates in the conceptual phase (AACE International, 2020).

From a practical perspective, this requires us to view feasibility planning and preliminary design as a dynamically coupled process, rather than isolated linear steps. Once the team artificially separates regulatory review, cost estimation, and design simulation into isolated workflows, they are often forced to bear high rework time costs due to delays in identifying core constraints (Miles et al., 2015).Table 2.3 provides a simple feasibility checklist that translates administrative requirements into evidence tasks. It also highlights typical bottlenecks and their consequences for schedule and cost (D.P.R. 380/2001, art. 10–12; AACE International, 2020).

Table 2.3. Administrative/Planning Feasibility Checklist (Traditional) + Typical bottlenecks + consequences for schedule/cost

Checklist item	Evidence needed	Main actors	Typical bottleneck	Schedule impact	Cost/feasibility impact
Identify permit triggers	Scope description; intervention type	Developer; designer; municipality	Scope unclears	Resubmission/review cycles	Redesign and sunk costs
Check conformity with plans/rules	Compliance drawings; statements	Designer; municipality	Interpretation disputes	Long review loops	Delay-driven value loss
Confirm infrastructure conditions	Servicing plans; commitments	Developer; utilities; municipality	Multi-agency coordination	Critical-path delays	Unbudgeted works
Check urban standards (e.g., distances)	Massing checks; sections	Designer; municipality	Non-compliant massing	Redesign iterations	Reduced buildable area
Manage engagement/negotiation	Meeting record; revision log	Developer; public bodies; stakeholders	Late objections	Uncertain timeline	Scope creep; mitigation costs
State cost estimate maturity	Estimate basis; assumptions	Cost consultant; developer	Overconfidence in early estimate	Late cost shock	Budget overruns; financing strain

Note. Permit and conformity logic follows D.P.R. 380/2001, arts. 10 and 12. Urban standards example reflects D.M. 1444/1968, art. 9. Estimate maturity follows AACE International (2020).

Mini summary

Planning feasibility and early design are tightly linked because permits and conformity checks set binding limits on what is buildable and on the project timeline. Approvals are often iterative, and delays can reshape both costs and market assumptions. Cost estimates are uncertain at concept stage, so teams need explicit accuracy statements to avoid overconfidence (D.P.R. 380/2001, art. 12; AACE International, 2020; Miles et al., 2015).

Common pain points in practice

- Approvals often sit on the critical path and can require multiple submissions, which adds time and consultant cost (Miles et al., 2015).
- Regulatory constraints (e.g., minimum distances) can invalidate early massing and force redesign (D.M. 1444/1968, art. 9).
- Infrastructure and urbanization obligations may appear late and create cost/schedule shocks (D.P.R. 380/2001, art. 12).
- Concept-stage cost estimates have wide accuracy ranges, but teams may still treat them as precise (AACE International, 2020).
- Coordination across agencies and disciplines increases documentation burden and slows iteration (Miles et al., 2015).

2.4 Limitations of the traditional workflow

The traditional workflow produces a structured sequence of outputs, but its weaknesses tend to compound across the market, design, and financial streams. The limits below are presented as an integrated synthesis so that the practical consequences for decision quality are clear even when individual tasks are performed correctly (Manganelli, 2015, p. 36; IVSC, 2024; RICS, 2024).

Evidence base and data limitations

Early-stage decisions often rely on incomplete or uneven evidence, especially for micro-locations where outcomes are most sensitive. When transaction data are scarce or delayed, analysts depend on proxies such as quotations or broad indicators, and those proxies may be carried forward without being clearly re-tested. This can bias revenue assumptions and make later comparisons unreliable because the evidence base is not consistent across iterations (Manganelli, 2015, p. 36).

Time and approvals uncertainty

Schedule uncertainty is a central limitation because time affects both market validity and discounted value. In practice, approvals and coordination requirements can create iterative cycles that extend timelines beyond early expectations. When timelines shift, market evidence can become stale and cash inflows arrive later than assumed, which reduces value and increases exposure to changing conditions (Manganelli, 2015, p. 75; Miles et al., 2015).

Cross-disciplinary coordination and iteration costs

Traditional workflows are often organized in silos, with market, design, planning, cost, and finance handled by different actors. Integration then happens late, after key assumptions have already been used in separate work products. When a design revision occurs, dependent inputs may not be updated consistently, leading to rework, duplicated effort, and misalignment between what is being tested and what is being designed (Manganelli, 2015, pp. 89–92; Miles et al., 2015).

Model fragility and sensitivity to key assumptions

Financial outputs can be highly sensitive to a small set of assumptions, especially timing, discount rate, and exit yield. Terminal value can drive the result: in many models, most of the final value comes from the assumed exit price, not from cash flows during development. As a result, a model can look numerically stable while still being conceptually weak if these assumptions are not supported by strong evidence (Manganelli, 2015, p. 104; Manganelli, 2015, pp. 133–134; RICS, 2024).

Documentation and governance (traceability and overrides)

The lack of model governance mechanisms is often a deep-seated cause of valuation instability. In practice, core assumptions are often highly fragmented, easily scattered across various spreadsheets, emails, or frequently iterated version documents. Furthermore, manual intervention often lacks standardized update logs, directly weakening data traceability and making it exceptionally difficult and tedious to review variations in results during iterations. It is precisely to mitigate this risk that professional standards emphasize strict constraints on the scope of evaluation, data boundaries, and reporting transparency. The core intention is to prevent users from becoming overly reliant on seemingly accurate but unverified outputs (IVSC, 2024; RICS, 2024).

Decision-making and information asymmetry

Simple thresholds such as hurdle rates can support quick screening and stakeholder communication, but they can also hide key drivers. If a project fails a threshold, teams may not learn whether the cause was timing, market assumptions, cost structure, or financing constraints. This can lead to repeated errors across projects and less informed negotiation about what changes would improve feasibility (Moorhead et al., 2024).

Mini summary

The traditional workflow remains widely used because it produces familiar decision artifacts and aligns with standard professional reporting. Its limitations arise when evidence is weak, timelines are uncertain, disciplines work in parallel without consistent updates, and models rely heavily on a few sensitive assumptions. In these conditions, documentation and governance become as important as technical methods, because poor traceability increases the risk of misinterpretation and repeated rework (Manganelli, 2015, p. 36; Manganelli, 2015, p. 104; IVSC, 2024; RICS, 2024).

3. Literature Review: AI and Digital Tools in Real-Estate Decision-Making

Early-stage redevelopment decisions are commonly made with incomplete information, changing design options, and uncertain approval timelines. In practice, teams rely on partial market evidence, evolving cost estimates, and spreadsheet-based feasibility models that can be difficult to audit and update. Digital tools and AI methods are increasingly used to reduce these frictions, but they can also introduce new risks related to reliability, governance, and accountability (National Institute of Standards and Technology [NIST], 2023).

This chapter reviews the main AI approaches and the digital foundations that shape planning, design, and feasibility work. It then critically discusses what these tools can improve, what remains difficult, and what conditions are needed for responsible use in an EU/Italy context (European Union, 2016, 2024).

3.1 Overview of AI developments relevant to planning, design and feasibility

In the built environment, it is often more useful to explain AI by the tasks it supports than by a long technical history. This review groups AI tools into five practical types: (1) rule or checklist tools, (2) prediction tools, (3) pattern-recognition tools for images or text, (4) generative tools that produce options, and (5) large language models that support document work (Goodfellow, Bengio, & Courville, 2016; NIST, 2023). These categories map to common concept-stage needs such as compliance screening, demand estimation, option exploration, and documentation.

Rule or checklist tools apply explicit rules and thresholds. In planning practice, this can look like a structured check against zoning limits. Their main limitation is brittleness: if rules change, or if a regulation is interpreted differently, the tool can mislead unless its logic is updated and reviewed (NIST, 2023).

Prediction tools estimate an outcome from historical data. Machine learning (ML) refers to models that learn patterns from data to make predictions, for example price, rent, absorption speed, or energy demand. These tools can support early screening, but they depend strongly on local data quality and can perform poorly when transferred to a different city or market segment (Rosen, 1974; El Jaouhari, Aouinti, & Oubbati, 2024).

Pattern-recognition tools extract information from images and text. For example, computer vision can classify land cover from satellite imagery, and natural language processing can help organize planning documents or property descriptions. The key limitation is context: regulations include exceptions and case-specific language, and image-based proxies rarely capture legal constraints and micro-location effects (NIST, 2023).

Generative tools produce candidate options. In concept design, this can include systems that propose massing variations under constraints or generate alternative layouts. The benefit is speed in exploring options, but the risk is false precision: outputs can look validated even when the assumptions behind them are weak or undocumented (NIST, 2023).

Large language models (LLMs) generate and transform text. An LLM is a model trained on large text datasets to predict and produce language. LLMs can support tasks such as summarizing documents, drafting reports, and translating between languages. However, they can produce fluent but incorrect statements (often called hallucinations), and their outputs can vary across runs, which raises governance concerns in professional settings (Ji et al., 2023; NIST, 2024).

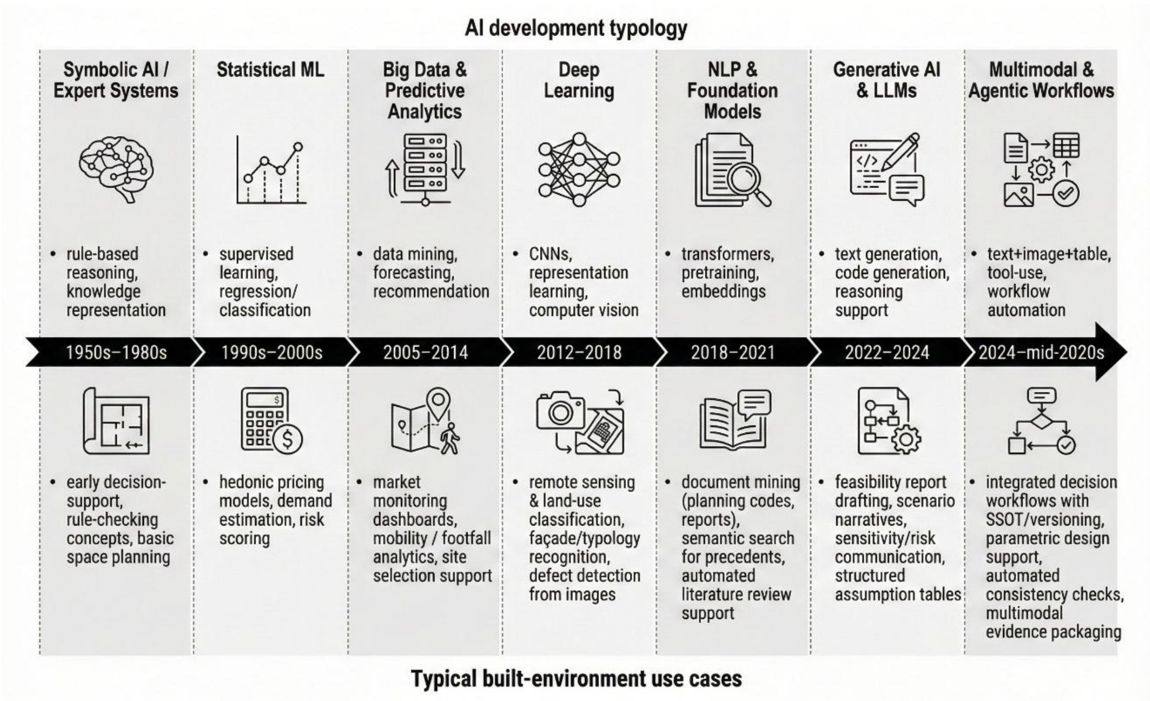

Figure 3.1. Timeline / typology of AI developments and typical built environment use cases.

The figure illustrates the five practical AI types described above. For each type it would show one built-environment use case (e.g., compliance checks, price/rent prediction, regulation parsing, option generation, document summarization) and one typical limitation (e.g., brittleness, transferability, noisy inputs, false precision, hallucinations) (NIST, 2023, 2024).

3.2 Applications of AI in architecture, planning and real estate

The literature reports many AI applications across the real-estate development workflow. For an architecture-led feasibility focus, it is useful to organize applications by early-stage decisions: market and demand judgement, valuation and risk framing, design and planning exploration, and communication across stakeholders. Across these areas, a key distinction is between what can be shown in research settings and what is deployable in real projects under data gaps, accountability requirements, and legal constraints (NIST, 2023; Royal Institution of Chartered Surveyors [RICS], 2024).

3.2.1 Market and demand analytics

Market analysis tools aim to estimate what can be built and what will be demanded at a given site. Common tasks include identifying comparable locations, estimating achievable rents or sale prices, and forecasting absorption. Automated valuation models (AVMs) are one important family: they are statistical or ML models that estimate property value from data. AVMs can perform well when detailed local datasets exist and when models capture location, building attributes, and time effects (Rosen, 1974; El Jaouhari et al., 2024).

In practice, market evidence is uneven. Transaction data can be delayed or proprietary, while listing data may reflect asking prices rather than final deals. This creates bias risks: a model trained mainly on listings may learn optimistic price levels. Valuation practice therefore stresses the need to cross-check evidence from multiple sources and to record which inputs are observed data versus assumptions (Manganelli, 2015, pp. 35–36).

Transferability is another constraint. A model trained in one city can perform poorly in another because the meaning of location variables changes with urban structure, mobility patterns, and regulation. Even within one city, micro-location effects are difficult to capture with standard datasets. For feasibility decisions, predictive accuracy alone is not enough; the model must be transparent about scope and uncertainty (NIST, 2023; RICS, 2024).

3.2.2 Valuation and risk assessment

Valuation and feasibility often rely on discounted cash flow (DCF) logic. DCF estimates value by forecasting cash inflows and outflows over time and discounting them to the present. In early stages, small changes in assumptions can produce large changes in outcomes, especially when the terminal value is a large share of the result (Manganelli, 2015, pp. 97–104; Moorhead, Armitage, & Skitmore, 2024).

Timing matters because development cash flows are not symmetric. A delay shifts revenues later, while many costs occur earlier. Even if total profit stays similar, later cash inflows are worth less once discounted. This is one reason why schedule uncertainty and approvals risk can dominate early feasibility judgements (Manganelli, 2015, pp. 97–100).

Professional valuation standards still expect transparency, appropriate evidence, and clear reporting of assumptions. Both IVS and the RICS Red Book emphasize scope definition, data adequacy, and documentation that allows the valuation to be understood and reviewed (International Valuation Standards Council [IVSC], 2024; RICS, 2024).

This creates a tension: complex ML models may predict well but can be hard to explain. Explainable AI methods (such as LIME and SHAP) can show which inputs are associated with a prediction, but they do not guarantee that the underlying data are unbiased or that relationships are causal (Lundberg & Lee, 2017; Ribeiro, Singh, & Guestrin, 2016). In valuation contexts, explainability must be paired with evidence quality and disciplined reporting (IVSC, 2024; RICS, 2024).

3.2.3 Design and planning decision support

In design and planning, AI is often used to explore options faster than manual iteration. Generative and optimization approaches can produce massing alternatives that satisfy constraints such as setbacks, maximum height, or target floor area. They can also run early proxy simulations for daylight, solar exposure, or basic energy demand. Their practical value is speed in exploring "what if" questions during concept development (NIST, 2023).

At the same time, early-stage simulations rely on simplified assumptions. If teams treat proxy outputs as final evidence, they risk locking in a concept that later fails more detailed checks. This is a typical false-precision problem and it can increase redesign costs (NIST, 2023; RICS, 2024).

Planning decisions also involve values, negotiation, and legal interpretation. AI can support evidence processing, but it cannot replace the institutional process that determines what is permitted and acceptable. For this reason, planning literature emphasizes accountability, bias risks, and the need for transparent decision rationales (NIST, 2023).

3.2.4 Participation, communication, and document work

Many redevelopment tasks are document-heavy: reading regulations, summarizing stakeholder feedback, and drafting reports. LLMs can help with these tasks because they can summarize text, translate, and generate structured drafts. In an architecture-led workflow, this can reduce time spent on repetitive writing and support clearer communication across disciplines (NIST, 2024).

However, LLM outputs require careful review. LLMs can produce confident statements that are not supported by sources, and they can mix incompatible assumptions. Scholars warn that fluency does not guarantee reliability and can hide systematic bias (Bender et al., 2021; Ji et al., 2023). Risk management guidance therefore emphasizes human oversight, documentation of prompts and inputs, and clear boundaries on where AI outputs are used (NIST, 2023, 2024).

In the EU context, data protection and accountability also matter. Personal data and sensitive project data must be handled carefully under GDPR. The EU AI Act adds risk-based obligations for certain AI systems, reinforcing the need for governance even when tools appear "low risk" in everyday use (European Union, 2016, 2024).

3.3 Digital and data-driven tools in feasibility and early-stage design

Many "AI problems" in real estate are data and workflow problems first. If information is scattered across drawings, emails, spreadsheets, and disconnected models, AI can amplify inconsistency rather than reduce it. For this reason, digital foundations such as BIM, GIS, and structured scenario management are often prerequisites for reliable AI use (International Organization for Standardization [ISO], 2018a, 2018b; Panko, 1998).

3.3.1 BIM, CDEs, and interoperability

Building Information Modelling (BIM) is a process for creating and managing structured information about a building or infrastructure asset. A Common Data Environment (CDE) is the shared digital space where project information is stored, versioned, and reviewed. ISO 19650 describes principles for information management using BIM, including naming conventions, responsibilities, and approval workflows (ISO, 2018a, 2018b).

For early-stage feasibility, BIM and CDE practices matter because they reduce rework and make assumptions traceable. Interoperability remains a challenge across software tools. The IFC standard is an open schema for sharing BIM data, but consistent exchange still requires modelling conventions and validation (ISO, 2024; building SMART International, n.d.).

3.3.2 GIS, urban datasets, remote sensing, and digital twins

Geographic Information Systems (GIS) manage spatial data such as parcels, land use, transport networks, and environmental constraints. In early redevelopment work, GIS supports site screening and context analysis, such as proximity to transit, flood risk screening, or accessibility indicators. Remote sensing adds image-based evidence, but it usually provides proxies rather than legal certainty (Boje, Guerriero, Kubicki, & Rezgui, 2020).

Urban digital twins are often described as digital representations that connect a place with data over time. Research stresses that digital twins require clear data definitions and governance. Without this, they risk becoming fragmented dashboards that do not support accountable decisions (Boje et al., 2020).

3.3.3 Parametric design and early-stage simulation

Parametric design uses explicit parameters (such as floor-to-floor height, setbacks, or core size) to generate geometry and enable rapid iteration. When connected to simulation, parametric models can provide early feedback on daylight, solar exposure, or basic energy demand. This supports option exploration, but results are only as good as the assumptions and boundary conditions provided (NIST, 2023).

A practical risk is that teams treat early proxies as fixed evidence. If constraints change, the model must be updated and rechecked. Without clear versioning, old results can persist in slides and feasibility models, creating assumption lock-in and decision confusion (ISO, 2018a, 2018b).

3.3.4 Financial modelling infrastructure and model risk

Financial models remain central in feasibility work, but they create model risk. Spreadsheet models can hide errors through manual overrides, copied formulas, and inconsistent scenario inputs. Research on spreadsheet errors shows that even experienced users make mistakes, and these errors are hard to detect without structured controls (Panko, 1998).

Digital tools can reduce these risks by improving scenario management, audit trails, and reproducibility. However, digitalization can also amplify false precision if teams rely on a clean interface while ignoring weak evidence behind key inputs such as the exit price or discount rate (Manganelli, 2015, pp. 97–104; Moorhead et al., 2024).

Table 3.1. Digital tool stack in early-stage feasibility: purpose → typical outputs → common pitfalls → governance needs.

Digital foundation / tool family	**Purpose (task)**	**Typical outputs**	**Common pitfalls**	**Governance needs**
BIM models	Concept geometry and quantities	Areas, volumes, schedules	Unclear modelling assumptions; inconsistent level of detail	Model versioning; documented modelling rules (ISO 19650)
CDE (ISO 19650)	Shared information management	Controlled revisions; approvals; audit trail	Teams bypass the CDE; parallel shadow files	Roles/responsibilities; naming; approval workflows
IFC / interoperability	Share models across software	Exchangeable BIM datasets	Loss of semantics; inconsistent mappings	Validation checks; agreed exchange requirements
GIS and urban datasets	Site/context screening	Maps, indicators, constraints layers	Outdated layers; scale mismatch; missing micro-location factors	Data provenance; update cycles; documented limits
Parametric + proxy simulation	Fast option generation and performance feedback	Massing options; daylight/solar proxies	False precision; stale results after design changes	Scenario labelling; rerun triggers; assumption logs
Financial models + scenario management	Feasibility evaluation and comparability	DCF, NPV/IRR, sensitivity outputs	Hidden overrides; inconsistent inputs; timing errors	Single source of truth; audit trail; version control

Note. External foundations and governance expectations align with ISO 19650 information management principles (ISO, 2018a, 2018b) and professional valuation documentation norms (IVSC, 2024; RICS, 2024).

3.4 SWOT analysis of available AI and digital tools

The SWOT analysis presented in this section is the author's evaluation based on controlled tool testing, access checks, and workflow trials conducted in December 2025. Tools were assessed against practical usability constraints relevant to an EU/Italy student redevelopment context: affordability and access, availability of local data, interoperability with common formats, learning curve, and the ability to document and reproduce outputs. The evaluation does not claim to measure absolute tool performance; it focuses on deploy ability under realistic constraints and accountability expectations (NIST, 2023; RICS, 2024).

Table 3.2. Evaluation criteria used in the SWOT: criteria → why it matters → supporting literature/standard.

Criterion	Why it matters in feasibility/design decisions	Supporting literature/standard (examples)
Evidence transparency	Decisions require a clear link from outputs to data and assumptions.	IVSC (2024); RICS (2024)
Auditability and traceability	Teams must explain changes, versions, and the basis for conclusions.	ISO 19650 (ISO, 2018a, 2018b); NIST (2023)
Interoperability	Workflows cross many tools; poor exchange increases rework and errors.	IFC / ISO 16739-1 (ISO, 2024); buildingSMART International
Local-data availability	Predictions and benchmarks fail when local market and rules are missing.	El Jaouhari et al. (2024); Manganelli (2015)
Privacy and compliance	Project and personal data must be handled under EU legal requirements.	GDPR (European Union, 2016); EU AI Act (European Union, 2024)
Reliability and error modes	AI may fail silently; risk controls are needed for safe use.	NIST (2023, 2024); Ji et al. (2023)
Cost and access	High fees or restricted access can block adoption in practice.	Tagliaro & Zanni (2020); Ullah et al. (2021)
Learning curve and workflow fit	If a tool is hard to learn or too narrow, it will not integrate.	Ullah et al. (2021); ISO 19650 principles

3.4.1 SWOT Analysis of Vertical AI and Digital Tools

Vertical tools focus on a specific task, such as early massing, market benchmarking, cost estimation, or lease modelling. In the author's December 2025 evaluation, the main barrier was often not the lack of features, but deploy ability: cost, access restrictions, missing or weak Italy-relevant datasets, and limited interoperability. Even when a tool worked well for one step, integration across the full workflow was difficult due to format mismatches and duplicated assumptions (Tagliaro & Zanni, 2020; Ullah et al., 2021).

Table 3.3. Practical usability filter scoring table (Author's evaluation, Dec 2025).

Stage	Tool	Type	Cost	Italy data	Integ-ration	Onbo-arding	Output fidelity	GDPR /Gov risk
Cross-stage	Large Language Models (LLMs)	AI (general)	4	3	4	4	3	3
1	CoreLogic	AI-enabled data platform	1	1	3	2	3	3
1	HouseCanary	AI analytics platform	1	1	2	2	3	3
1	PriceHubble	AI valuation/insights	2	3	3	3	3	3
1	Casafari	AI-assisted data aggregation	2	3	3	3	3	3
1	CoStar	Data platform (commercial RE)	1	1	3	2	3	3
2	Autodesk Forma	AI-assisted site/massing	4	3	4	4	4	3
2	Archistar	AI-assisted feasibility/design	2	1	2	2	3	3
2	TestFit	AI-assisted layout/feasibility	2	2	3	3	4	3
2	Maket.ai	Generative design (concept)	2	2	2	3	2	3
2	CityBldr	Site selection/feasibility	1	1	2	2	2	3
3	Deepblocks	AI feasibility platform	2	1	2	2	3	3
3	Argus Enterprise	Baseline (non-AI)	1	3	4	2	4	4
3	HouseCanary (forecast module)	AI module	1	1	2	2	3	3
3	PriceHubble (finance module)	AI module	2	3	3	3	3	3
3	EstateMaster / Excel model	Baseline (non-AI)	5	5	5	3	5	5
4	Delve (Sidewalk Labs)	Reference concept / limited access	U	U	U	U	3	U
4	UrbanFootprint	Planning decision-support platform	1	2	3	1	3	3
4	Archistar +Deepblocks	Integration pattern	1	1	2	1	3	3

Source: Author's tool testing and comparative evaluation (Dec 2025). Scoring scale: 1 = poor / not usable in this context; 3 = usable with clear constraints; 5 = strong / practical. "U" = unknown (not verified or not testable within the project's access limits). Vendor documentation was consulted where available to check access/pricing/features (see References). Scores are a deploy ability screen (access, data coverage, workflow fit), not an absolute performance benchmark.

Table 3.4. SWOT table (Stage 1) — Market / data and benchmarking tools (Author's evaluation, Dec 2025).

Tool	Strengths	Weaknesses	Opportunities	Threats
CoreLogic	A global-leading real-estate data and analytics provider with broad coverage; offers detailed property history, valuations, and market-trend data; a mature product widely used in North America (the exact user scale varies by source and definition).	Focused on data provision rather than decision recommendations; coverage concentrates on North America and Australia, with limited local data support for Italy; expensive data services with no clear free tier or education discount, making it unfriendly to individual/student users.	Can increase insight value by pairing with AI analytics (e.g., using machine learning to detect market patterns and provide early risk signals); its large customer base suggests strong potential if the platform expands deeper into European markets.	Competition from emerging PropTech data platforms; data-privacy rules can restrict cross-border data sharing; in markets with weak localization (e.g., Italy), practical usefulness is limited if robust local datasets cannot be provided.
HouseCanary	An AI-driven real-estate market analytics platform that provides comprehensive data for decision-making; includes functions such as forecasting prices, rental returns, and risk scores; User feedback suggests it can reduce the time required for comparable-market analysis (CMA) tasks.	The interface and advanced analytics can be complex; new users may need training to use it effectively; mainly oriented to the U.S. market, with no local support for Italy; subscription pricing (varies by plan/region and should be verified), typically without a long-term free version.	As AI adoption in valuation grows, HouseCanary could expand to more markets and integrate with broader investment decision workflows; its predictive outputs could support feasibility assessment and risk management for development projects.	Geographic limitation: without entering European markets, it has little direct applicability in Italy; faces competition from domestic data platforms such as Zillow; reliance on predictive models also creates the risk of underperformance under abnormal market conditions.
PriceHubble	A leading European AI valuation and insights platform, using big data and machine learning to forecast residential prices and provide market insights; operates across multiple countries (e.g., France, Germany, Switzerland) and	Primarily focused on residential valuation, with limited support for full development feasibility (no integrated planning/design and financial modules); data quality depends on public datasets and partners and may be incomplete in some regions; largely B2B, so individuals	Potential to enter new markets including Italy and fill gaps in intelligent valuation tools; its value and rent estimates can support market demand and pricing assumptions for the Spina project; may also add more "AI agent"	Established national valuation systems (e.g., Italy's Agenzia delle Entrate OMI data) may substitute for some use cases; competition from European peers such as RealAdvisor and Casafari; if the market does not trust AI valuation accuracy, adoption may be constrained.

Tool	Strengths	Weaknesses	Opportunities	Threats
	supports banks with customer-facing tools and credit-risk management; emphasizes model transparency and explainability.	may need access via an institution.	features to assist real-estate decision-making.	
Casafari	A pan-European real-estate big-data platform collecting listing and transaction information across multiple countries including Italy; uses AI to merge and deduplicate large datasets, providing market analysis, valuation support, and investment leads; strong at cross-source integration and identifying undervalued assets and opportunities.	Mainly serves brokers and institutional investors, creating a higher entry barrier for individual users; relies largely on listings and public sources and may miss off-market transactions; lacks dedicated concept-design or financial-analysis modules and therefore needs to be paired with other tools.	As PropTech grows across Europe, Casafari could become a standard data layer feeding other AI analytics tools; its market-monitoring functions can support early location screening and, for projects like Spina , help identify candidate sites and price trends.	Restrictions on data sources (e.g., anti-scraping policies or GDPR enforcement) could reduce completeness; competition from local data providers (e.g., Italy's Nomisma); if it cannot deliver higher-value analytics, it may be displaced by more comprehensive platforms.
CoStar	One of the world's largest commercial real-estate (CRE) data and analytics platforms, with strong authority in CRE; covers sales/lease transactions, inventory, rent levels, and future supply, supporting market research (platform-scale metrics often come from public marketing claims and vary by source/time and should be verified); mature firm with high data accuracy.	Expensive subscriptions with strict enterprise licensing, making access difficult for individuals and education; mainly active in North America and parts of Europe (e.g., the UK), with limited coverage in Italy; focused on data query/analysis and offers limited generative design or automated feasibility-judgement functions.	If CoStar expands coverage in continental Europe, it could provide strong market-research support for projects like Spina ; rich datasets could be combined with AI decision algorithms for smarter site selection and investment analytics; acquisitions may introduce new AI features and improve user experience.	Emerging platforms (e.g., Crexi, Reonomy) challenge incumbents with more open or lower-cost models; if CoStar does not integrate AI quickly, it may fall behind intelligent decision-support trends; for Spina , insufficient local data could create biased conclusions if relied upon heavily.

Source: Author's tool testing and comparative evaluation (Dec 2025).

Table 3.5. SWOT table (Stage 2) — Planning / early design tools (Author's evaluation, Dec 2025).

Tool	Strengths	Weaknesses	Opportunities	Threats
Autodesk Forma	A leading early-stage AI design platform that enables rapid concept testing and site optimization; includes built-in environmental simulations (energy, daylight, wind, mobility, etc.) and regulatory checks to improve concept robustness; integrates with Autodesk's BIM toolset for downstream detailed design.	Steep learning curve; advanced use takes time to master; subscription costs can be high (often unaffordable for small studios or individuals); for highly non-standard design intents, automated generation remains limited and still requires designer-led refinement.	Autodesk offers a free student license, enabling education users to use the platform at no cost for one year; continuous updates may add further functions (e.g., automated financial feasibility checks), moving toward an end-to-end concept tool; for sustainability-focused projects like Spina , Forma's environmental analytics can support green optimization strategies.	Growing competition from similar platforms (e.g., Archistar) may erode market share; overly complex or "black-box" optimization can reduce decision-makers' sense of control; reliance on cloud workflows and high-quality base data means performance may degrade where mapping/parameter data are weak.
Archistar	An integrated AI development platform combining site selection, regulatory analysis, generative design, and basic financial feasibility; generates large numbers of massing/concept options within zoning/height limits and checks compliance in real time (e.g., height and daylight overshadowing); offers a free basic tier, lowering the barrier for small teams and student users.	Optimized mainly for multifamily/apartment projects and less suitable for specialized buildings or complex urban design; aesthetics and functional quality of auto-generated options can be uneven and require designer review; detailed datasets and templates are primarily Australia-focused (integrating major Australian data sources), so users outside Australia (e.g., Italy) must input regulations and data manually, increasing effort.	Potential to expand globally by transferring its Australian experience to Europe (if local planning data can be integrated); the eCheck digital-approval feature (AI-assisted plan checking) could create additional value; for Spina , Archistar could support early site-capacity tests and concept comparisons, and its finance module can help compare development returns across options.	Overlapping functionality with other tools (e.g., Autodesk Forma) means large vendors could compress its niche; without strong European planning datasets, user effort and cost increase while reliability falls; generative outputs depend heavily on accurate inputs—errors in regulation/data can misdirect design decisions and waste time.
TestFit	Fast concept generation and testing tool, sometimes described as a "feasibility calculator". Strong for standardized modules such as apartments/multifamily and parking layouts, enabling rapid parameter-driven iteration. Users report	Limited scope: mainly suited to typical residential and parking typologies, with weaker support for complex mixed-use or urban-design cases; relies on user-provided local parameters (e.g., costs, sales values) to run economic tests	Recent product updates expanded generative-design capabilities and the range of supported building types. In education settings, classroom licenses and student trials can make it a useful learning tool for	Large vendors (e.g., Autodesk) may integrate similar real-time generation into BIM platforms, creating direct competition; local regulations and market parameters still require manual updating and are not automatically

Tool	Strengths	Weaknesses	Opportunities	Threats
	that real-time scenario testing can support internal communication and cross-team coordination.	and therefore requires domain expertise; while a beta Revit plug-in exists, integration with mainstream BIM tools is still developing.	early-stage feasibility exercises.	refreshed by the system; for a specific site like Spina , TestFit's built-in U.S.-centric standards (e.g., parking ratios, unit sizes) may be inappropriate and require customization—otherwise outputs may be unrealistic.
Maket.ai	An emerging AI architectural concept platform focused on automatic generation of residential floor plans and 3D layouts; can produce multiple unit-plan options in one click and support style exploration (via prompt-based variation); offers a "regulation assistant" intended to help check zoning rules (under development); can integrate with common design software such as Revit and SketchUp, enabling export for further development.	Currently focused on residential buildings and is less general for complex or commercial projects; generated layouts can be formulaic and typically require designer revision to achieve project-specific quality; subscription pricing applies after trial use, with no clearly stated education discount.	With more training data, layout quality may improve and extend to more building types; once the regulation-assistant function matures, it could add practical value at concept stage; if it enters European markets, Maket.ai could provide fast plan-layout references for residential-led redevelopment projects like Spina , reducing early design time.	Intense competition from similar platforms (e.g., Architectures, Planalogic) means user loyalty is not yet stable; legal risks around authorship/copyright and accountability for AI-generated designs remain unclear; for highly creative urban design, over-reliance on such tools may suppress innovation.
CityBldr	An AI platform focused on "highest and best use" prediction: it scans parcels, assesses development potential, and recommends higher-value development options; models are trained with engineering and appraisal expertise to identify undervalued sites and redevelopment opportunities; can connect sellers and buyers, supporting a one-stop workflow from discovery to transaction execution.	Currently concentrated on U.S. cities with no direct coverage for Italy; more of a land-screening and valuation tool than a building-design generator; limited transparency of data and modelling means results require professional review, and local factors (e.g., heritage constraints, community sentiment) may not be fully captured.	If expanded to Europe, CityBldr could help public bodies and developers identify "sleeping assets" and discover developable sites in cities such as Turin, including sites similar to Spina ; by suggesting likely best-use mixes, it could provide an initial direction (e.g., office vs residential) before handing off to design tools; it also has potential value for urban-regeneration acquisition and negotiation contexts.	Planning changes or abrupt market shifts can cause prediction failures (e.g., demand changes after a shock); local professionals' judgement remains a strong competitor—if AI recommendations conflict with practice intuition, they may be rejected; differing legal/policy environments require retraining and localization, which can be costly and difficult.

Source: Author's tool testing and comparative evaluation (Dec 2025).

Table 3.6. SWOT table (Stage 3) — Financial feasibility and risk tools (Author's evaluation, Dec 2025).

Tool	Strengths	Weaknesses	Opportunities	Threats
Deepblocks	An AI platform that combines early design exploration with simplified financial testing in a single cloud workspace. It can support fast iteration between "what can I build?" and "what might it be worth?", which is helpful in concept-stage screening.	Limited city coverage: full data integration currently exists only for selected U.S. cities; as a relatively young company (founded in 2016), functions are evolving quickly, and stability/support may be uneven; for non-covered cities (e.g., in Italy), users must import data themselves, increasing effort and reducing reliability.	Plans to expand to more regions (the founders indicate expansion beyond the U.S., potentially into Latin America); may add modules such as climate-risk analysis, widening use cases; if relevant base data can be assembled for Spina , Deepblocks could link zoning rules to financial outputs and allow real-time testing of alternative schemes, producing comparable reports that support option selection and demonstrate AI value.	Not yet established in Europe, limiting short-term usefulness for the Spina case; large PropTech incumbents or local firms may develop similar integrated tools, reducing differentiation; heavy reliance on public data means that delayed updates or errors in data sources could mislead financial outputs and decisions.
Argus Enterprise (baseline / non-AI finance software)	A long-standing benchmark for real-estate financial analysis, widely used by large international developers and valuers; supports detailed cash-flow modelling, hold/exit analysis, and sensitivity testing; mature and trusted by investment institutions.	Not AI-driven: does not provide automated forecasts and relies on analyst-specified assumptions; complex interface/workflows with a high learning burden for new users; expensive licensing, generally unaffordable for individuals.	When connected to AI data sources, Argus models could incorporate improved forecasts (e.g., AI-estimated occupancy or price-growth assumptions); the vendor is also exploring machine-learning applications in portfolio analytics, which may lead to "smart recommendation" features.	Newer AI platforms (e.g., Deepblocks) offer more user-friendly interfaces and automation, attracting some users; Argus's relatively closed ecosystem can limit integration with modern data workflows—without innovation it may be displaced; for students, access is extremely limited and typically depends on supervisors or corporate partnerships, constraining academic use.
HouseCanary (analytics functions)	HouseCanary's predictive analytics can also support the finance stage, providing forecasts of future prices and rents and market risk indices; these outputs can serve as key inputs for development revenues and exit values; its	Still largely confined to U.S. datasets, so it cannot directly forecast Italian market indicators; focuses on pricing individual residential assets rather than full development returns (including development costs and construction schedules); needs to	If expanded into Europe, its machine-learning models could help forecast Italian residential price trends and support scenario planning (upside/downside market cases); risk indices could support area screening and portfolio allocation (selecting lower-risk, higher-growth locations).	Limited geographic applicability reduces competitiveness relative to local data experts for Spina ; forecasts are model outputs and may not reflect "black swan" events (e.g., sudden policy changes) in time, so over-reliance is risky; Europe also has bank/institution in-house models, so

Tool	Strengths	Weaknesses	Opportunities	Threats
	automated comparable-selection methods can help position pricing more accurately (reducing noise from poor-quality comparables).	be combined with a separate cost/cash-flow model.		market entry would require trust building.
PriceHubble (finance module)	As a valuation tool, PriceHubble's value and rent estimates can inform revenue assumptions for development projects; credit-risk indicators developed for banks may also help developers understand financing risk premia; its emphasis on "explainable AI" supports interpretability and can improve credibility in financial discussions.	Does not directly model development costs or full cash flows and therefore must be paired with a feasibility model; in some countries it has limited data for office/commercial assets—if Spina includes significant non-residential uses, outputs may be restricted to the residential component; integration into feasibility spreadsheets may require custom workflows.	Opportunity to develop developer-oriented products (e.g., ROI reference ranges by asset type based on macro data and AI forecasts); for investment reports that require strong data backing, PriceHubble outputs can strengthen credibility and become a communication "highlight" in discussions with investors.	Banks and valuers remain cautious about AI-based appraisal; if key financial decisions rely heavily on AI outputs, they may be treated conservatively; competitors (e.g., RealAdvisor) may launch similar services and erode market share; under abnormal market conditions not represented in training data, miscalibrated outputs could materially mislead financial decisions.
EstateMaster / Excel models (baseline / non-AI approach)	Many European developers still use Excel or traditional tools such as EstateMaster for feasibility models; high flexibility enables full adjustment to local experience, parameters, and formulas; transparency is strong—teams can inspect each calculation step, aiding audit and review.	Manual modelling is time-consuming and error-prone and makes rapid iteration across many options difficult; without ML support, external market forecasts are not generated automatically and assumptions must be input manually; advanced uncertainty analyses (e.g., Monte Carlo simulation) are cumbersome to implement.	Can be combined with AI tools—for example, using Python interfaces to pull AI-based forecasts (e.g., from PriceHubble) into Excel models, enabling a hybrid "AI + analyst" approach; in academic work, custom ML models can also be trained to replace some assumptions and enhance traditional spreadsheets.	Efficiency and "intelligence" are far below new AI-native platforms, creating disadvantage in fast, competitive decision contexts; over-reliance on individual judgement can embed human bias; for students, using only traditional methods may under-demonstrate AI value and miss opportunities to validate assumptions quickly with new tools.

Source: Author's tool testing and comparative evaluation (Dec 2025).

Table 3.7. SWOT table (Stage 4) — Integrated / combined approaches (Author's evaluation, Dec 2025).

Tool	Strengths	Weaknesses	Opportunities	Threats
Delve (Sidewalk Labs)	A generative urban design decision-support tool that can produce many massing/layout options under defined constraints. Useful for exploring "what-if" alternatives quickly at district scale and for communicating trade-offs to stakeholders.	As a Sidewalk Labs product, Delve currently has limited public access (after organizational changes it is no longer broadly offered as a standalone product); requires detailed infrastructure and cost models, creating a high barrier for smaller projects with incomplete data; outputs are complex and often require expert interpretation, and the interface/analysis style is unfamiliar to many developers.	Delve illustrates the long-term potential of AI for integrated decision-making—if similar technologies become commercially available, teams could test multiple objectives (financial, social, environmental) within one platform; for projects like Spina that must balance economic viability with urban-quality goals, multi-objective optimization can help make trade-offs explicit and quantify why one option outperforms another.	Because Delve is not fully open, users can mainly learn from its concept rather than deploy it; if no mature comparable products emerge, developers may continue relying on manual decision-making and miss efficiency/innovation gains; if new platforms emerge, Delve's approach may be surpassed or replaced; moreover, if multi-objective AI outputs conflict with professional judgement, stakeholders may question legitimacy and resist adoption.
Archistar / Deepblocks (integrated end-to-end use)	Platforms such as Archistar and Deepblocks bundle multiple functions and can support a near end-to-end workflow. For example, Archistar can generate concept schemes and run zoning/compliance checks, while Deepblocks links 3D massing to pro-forma outputs; parameter changes allow users to see development indicators update in real time, enabling interactive "scheme testing". Such all-in-one	Although integrated, each sub-function may be shallower than specialized tools (e.g., simplified financial testing). For complex projects, a single platform rarely covers all analyses at the required depth (e.g., detailed cash flows, detailed design, and stakeholder requirements). Current geographic data limitations remain: Archistar is Australia-oriented and Deepblocks is U.S.-oriented, so	If basic datasets relevant to Spina can be assembled, an integrated platform could be used to iterate repeatedly across options—adjusting buildable envelope and program mix and immediately reviewing feasibility impacts. This is similar to searching for an "optimal" scheme in an AI sandbox and can produce persuasive comparative reports. Over time, these platforms may also integrate more optimization algorithms, further improving scheme recommendations.	Relying on one platform for the whole workflow creates concentration risk: if one module fails, the entire process can stall. Integrated platforms may also struggle to keep pace with best-in-class innovations (e.g., new financial modelling approaches or design algorithms), reducing the value of "all-in-one". For the Spina team, forcing a single-platform workflow could reduce insight in some dimensions, so tool selection should remain flexible and task-driven.

	tools reduce cross-platform import/export friction and can improve efficiency.	both require substantial data preparation for practical use in Italy.		
UrbanFootprint	A data-driven urban planning decision-support platform that integrates many spatial datasets (population, land use, mobility, environment, etc.) for scenario simulation and impact assessment; using AI and big data, it can quickly estimate how alternative development options affect transport, carbon emissions, infrastructure demand, and other outcomes, helping decision-makers consider social and environmental impacts alongside economics.	More oriented to macro-scale planning support and does not directly output micro-scale design or project-level financial analysis for a single site; primarily serves governments and planning agencies and can be complex to operate without a specialist team; Italian localization requires appropriate datasets—if data are missing, simulations become unreliable.	As sustainability requirements rise, tools like UrbanFootprint can add value in the comprehensive feasibility stage by integrating "externalities" (e.g., carbon, mobility, equity) into option evaluation. For projects that set carbon-reduction or accessibility goals, such tools can model trade-offs and support selecting an option that may have slightly lower ROI but better climate performance and policy alignment.	If developers focus only on financial returns, they may ignore these analyses and miss important dimensions in integrated appraisal; public-sector decision models may also feel redundant in private development contexts, limiting adoption; data privacy and openness constraints can further restrict practical usefulness in some countries.

Source: Author's tool testing and comparative evaluation (Dec 2025).

3.4.2 SWOT Analysis of Large Language Models

LLMs are not design engines, but they can support many text- and knowledge-heavy tasks in redevelopment work. In the author's December 2025 evaluation, LLMs were most useful for summarizing documents, generating structured checklists, drafting reports, and translating between languages. They were less reliable for tasks that require verified local facts, exact citations, or quantitative modelling without a controlled link to datasets and assumptions (Bender et al., 2021; Ji et al., 2023; NIST, 2024).

For feasibility and valuation contexts, a practical principle follows: LLMs can help structure and communicate reasoning, but they should not be treated as primary evidence for market values, legal compliance, or risk ratings. Professional standards expect key assumptions and evidence to be documented, reviewable, and appropriate to the valuation purpose (IVSC, 2024; RICS, 2024).

Table 3.8. SWOT of large language models for urban redevelopment decision-making (Author's evaluation, Dec 2025).

Model	Strengths	Weaknesses	Opportunities	Threats
ChatGPT 5.2	Strong general-purpose reasoning and structured writing for feasibility narratives, assumption logs, checklists, and reporting. Good support for iterative Q&A and synthesizing long documents (when sources are provided).	Not a geometry engine: cannot replace BIM/massing tools for compliant modelling. Outputs can contain errors or unsupported claims; requires human verification and clear source control. Cost, usage limits, and available features depend on plan/API and can change over time.	Combine with CAD/BIM and parametric tools to reduce time spent on document-heavy tasks (regulation summaries, meeting notes, option comparison). Use as a "glue layer" to keep assumptions consistent across market, design, and finance documents.	Over-trust risk: fluent text may be accepted as evidence without checking. Platform updates ("model drift") can change behavior between project iterations, affecting reproducibility. Privacy/compliance constraints limit what project data can be shared externally in some settings.
Gemini 3.0	Native multimodality (text + images) can support quick reading of diagrams, maps, and visual material. When paired with Nano Banana Pro (Gemini image-generation model), it can quickly produce and edit infographics, site-context visuals, and simple illustrative figures for reports and stakeholder communication. Strong ecosystem integration for users already working in Google tools.	Workflow can be "platform-dependent": the most convenient features often rely on the Google ecosystem. Data handling and organizational policies may restrict use for sensitive project information. Advanced multimodal/image features (including Nano Banana Pro) may be tied to premium tiers or restricted availability, limiting use in student or small-team contexts.	Potential for tighter links to web-scale information and geospatial datasets, supporting early context scanning and reporting. A practical path is a combined 'text + visuals' pipeline, where narrative drafts and figure mock-ups are created quickly, then manually checked, labelled, and documented for traceability. Useful for teams already using Google Cloud / Workspace.	Regulatory and governance scrutiny (privacy, accountability) can constrain adoption in public-sector contexts. Generated images can be misread as verified evidence unless clearly labelled and version controlled. Rapid model competition may reduce differentiation and increase switching uncertainty.
Claude Opus 4.5	Often valued for careful long-form writing and summarization of complex documents. Strong "safe and controlled" positioning may fit organizations emphasizing	Access and pricing are plan-dependent and can be a barrier for intensive use. Tooling and multimodal capabilities vary by product tier and may require workarounds for visual tasks.	Useful for drafting and refining policy-sensitive or formal documents (e.g., memos, stakeholder communications). Can complement design tools by improving clarity, consistency, and	If competing platforms offer similar writing quality plus stronger multimodality/tool integration, users may switch. As with all LLMs, hallucination and citation reliability remain risks

Model	Strengths	Weaknesses	Opportunities	Threats
	governance and compliance.		traceability of written outputs.	without strong source control.
Grok 4.1	Positioned for real-time awareness and fast response to current information streams. May be useful for monitoring news, sentiment, or rapidly changing context around a project.	Maturity and enterprise governance features are less established than long-standing competitors. Reliance on specific data sources/platform integrations can shape what it is good at (and what it misses).	Niche use in "monitoring" tasks: tracking policy debate, public discussion, or market news during long approval timelines. Could be combined with structured project datasets for more grounded reporting.	If data-source access changes, the tool's main advantage can weaken. A "real-time" posture can still mislead if sources are noisy or biased; verification remains necessary.
DeepSeek	Open-weight availability enables local or private deployment, supporting stronger control over data and model behavior. Can be adapted with local knowledge (rules, terminology, project templates) when technical capacity exists.	Requires technical setup and maintenance; user experience and support may be weaker than managed services. Model quality and safety controls vary across releases and deployments.	Opportunity for a "project-specific assistant" hosted in a controlled environment (e.g., within a university or firm). Useful where privacy constraints block sending data to external services.	Rapid advances by proprietary models can reduce the perceived benefit of self-hosting. Misconfiguration or weak governance can introduce security, reliability, and compliance risks.
Qwen3 Max	Large-scale model family with broad multilingual capability; potentially useful for bilingual documentation and cross-language evidence work. Can support structured writing, summarization, and translation tasks at scale.	Access, deployment options, and governance may depend on platform availability and organizational policies. For some EU/Italy contexts, procurement, data residency, or geopolitical concerns may limit use.	Potential fit for organizations already using the related cloud ecosystem, especially for translation-heavy workflows. Could support cross-border projects where multilingual synthesis is a frequent need.	Policy and procurement restrictions may limit adoption in Europe. As with other LLMs, without strong source control the outputs can still be confidently wrong.

Source: Author's tool testing and comparative evaluation (Dec 2025).

Note. The LLM SWOT summarizes the author's December 2025 comparative evaluation. Where specific product features are mentioned, they are treated as indicative and may change over time; the table therefore emphasizes qualitative implications for early-stage feasibility work. Vendor documentation and official release notes were consulted where available (see References).

3.5 Research gaps

Despite rapid research progress, several gaps remain between AI capabilities described in the literature and what can be used responsibly in early-stage redevelopment decisions, especially in an Italy/EU context.

First, there is an integration gap. Many studies and tools address a single task—price prediction, option generation, or document summarization—but feasibility decisions require an end-to-end chain from constraints and market evidence to design options and financial scenarios. Without integration, teams often maintain multiple inconsistent models, weakening comparability and accountability (ISO, 2018a, 2018b; Panko, 1998).

Second, there is a local-data gap. High-performing prediction models typically rely on detailed transaction datasets and stable attribute standards. In many contexts, micro-market evidence is fragmented, delayed, or proprietary, and regulatory texts are complex and language-specific. This limits transferability and increases the risk of spurious precision (El Jaouhari et al., 2024; Manganelli, 2015, pp. 35–36).

Third, there is an evaluation gap. Early-stage decisions have few shared benchmarks because outcomes unfold years later through approvals, design changes, and market cycles. This makes it difficult to compare tools, validate models, and learn systematically from past projects (NIST, 2023).

Fourth, there is a governance gap. Responsible use requires documentation, reproducibility, and clear responsibility for decisions. Generative models can be helpful, but their error modes (hallucinations, bias, and opaque reasoning) can conflict with professional expectations for transparent reporting and reviewability (Bender et al., 2021; IVSC, 2024; NIST, 2024; RICS, 2024).

These gaps motivate a structured framework that links tool selection to decision tasks, data availability, and governance requirements. The thesis therefore develops a practical workflow that is explicit about evidence, uncertainty, and traceability in early-stage redevelopment decisions.

Mini summary

This chapter reviewed how AI and digital tools are used in early-stage feasibility and design decisions. AI can support rapid screening, option exploration, and document work when tasks are well scoped (NIST, 2023, 2024). Prediction models can be useful, but they depend on local data quality and can fail when transferred across markets (El Jaouhari et al., 2024). Generative design tools can accelerate option exploration, but they can also create false precision if proxy results are over-trusted (NIST, 2023; RICS, 2024). LLMs can reduce document workload, yet they require strict human review due to hallucination and reproducibility limits (Bender et al., 2021; Ji et al., 2023). Across tools, digital foundations—information management, interoperability, GIS context data, and scenario governance—are prerequisites for reliable use (ISO, 2018a, 2018b). Overall, both literature and the author's testing highlight that deployability constraints (cost, data access, and integration) often matter as much as technical capability (Ullah et al., 2021).

Transition

To move from literature to application, the thesis develops a workflow that connects constraints, market evidence, design options, and financial scenarios through a controlled set of assumptions. The aim is to improve transparency and reduce iteration costs while keeping accountability for decisions with the human project team (NIST, 2023; RICS, 2024).

4. Integration of AI into the Real-Estate Decision-Making

This chapter operationalizes the thesis arguments from Chapters 2–3 into a practical, auditable methodology for integrating AI into early-stage real-estate decision-making in an EU/Italy context. The method is developed and tested using the Turin urban regeneration case study (SPINA 3 – Corso Principe Oddone), while the detailed site and design outcomes are reserved for Chapter 5. Here, the focus is on (i) workflow design, (ii) evidence packaging and upload discipline, and (iii) the selection of a large language model (LLM) through controlled stress testing and rubric-based scoring. Throughout, the goal is not automation, but decision support with traceability, version control, and explicit responsibility boundaries (ISO, 2018a, 2018b; NIST, 2023; IVSC, 2024; RICS, 2024).

4.1 Basic Process

The proposed methodology is structured as an end-to-end feasibility workflow in which each task is assigned to the tool type most capable of producing reliable and reviewable outputs. In line with the limitations discussed in Chapters 2–3, the workflow separates (a) evidence synthesis and narrative reasoning (LLM), (b) geometry- and rule-dependent feasibility checks (parametric / simulation platform), and (c) visual communication tasks that benefit from multimodal generation and rapid iteration (image model), while retaining a final human review as the accountability gate (NIST, 2023; IVSC, 2024; RICS, 2024).

Operationally, the workflow proceeds as follows: (1) consolidate and upload the project evidence package; (2) generate a single-run real-estate decision report in Deep Research mode (the artefact used for stress testing and scoring in Section 4.3); (3) extract a single-source of truth (SSOT) from the report into a structured assumption log for downstream modelling; (4) test technical feasibility through Autodesk Forma, focusing on massing, modelling, and early environmental constraints; (5) update the assumption log and financial/risk outputs through the LLM using controlled, versioned spreadsheet outputs; (6) synthesize a final feasibility report; (7) improve charts and figure quality using an image model; and (8) perform a final human review and deliver the report package with explicit version identifiers.

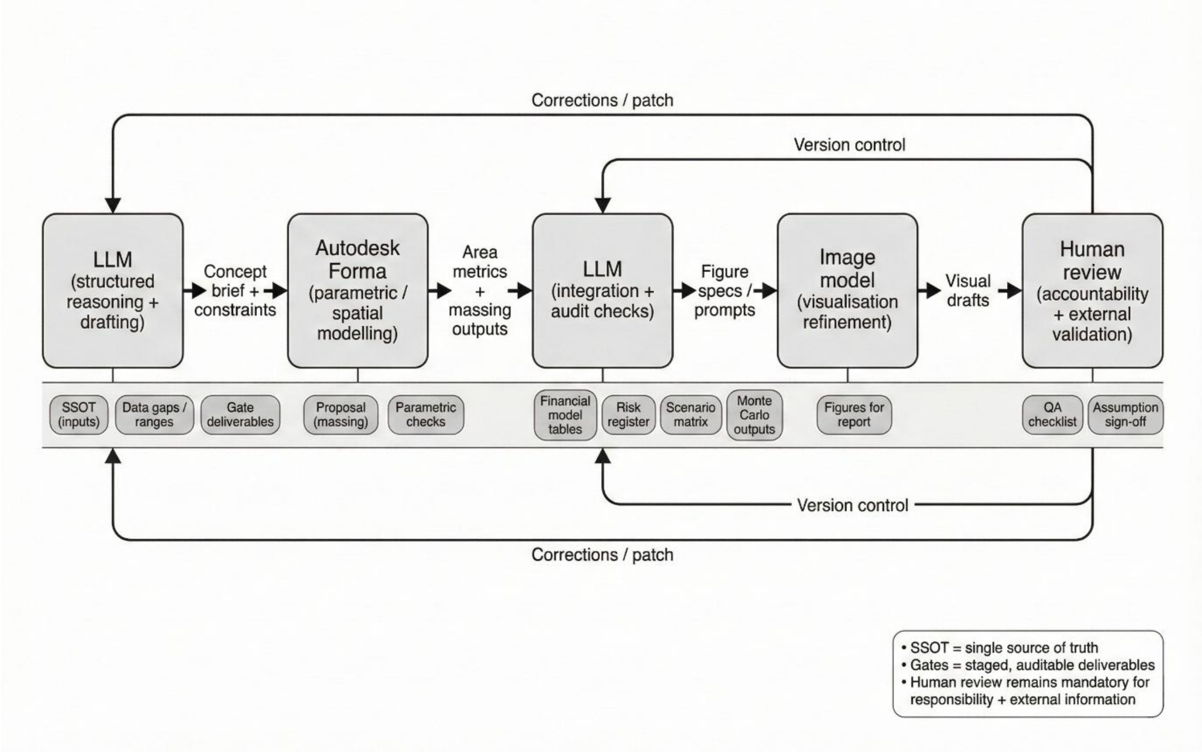

Figure 4.1. AI-integrated feasibility workflow for early-stage real-estate decision-making (methodological diagram).

The rationale for this division of labor is directly linked to known failure modes. LLMs are effective at summarizing heterogeneous documents, extracting qualitative and quantitative evidence, and producing decision-oriented narratives; however, they can omit hard constraints, overstate certainty, or introduce unsupported assumptions if evidence is incomplete or poorly organized (NIST, 2023). Parametric platforms, by contrast, make constraints explicit through geometry, rule sets, and measurable outputs, supporting verification of spatial feasibility and early performance checks that cannot be safely inferred from text alone. Finally, visual outputs in early feasibility reports often remain placeholders when produced purely as text; a dedicated image model can support the rapid production and refinement of charts and diagrams while the underlying numbers remain controlled in versioned spreadsheets (ISO, 2018a, 2018b; Panko, 1998).

4.2 Data Collection and Upload

Because LLM-based feasibility reasoning is only as reliable as the evidence it receives, data preparation is treated as a methodological step rather than an administrative task. In the Turin case study, the evidence package was constructed to cover (i) project and regulatory constraints, (ii) construction cost and timing benchmarks, (iii) socio-economic demand proxies, (iv) market price and rent evidence, (v) demographic structure and projections, and (vi) spatial-environmental context materials. The checklist below lists the datasets used in this chapter's workflow; Chapter 5 applies these inputs to the case-study design and feasibility outcomes.

4.2.1 Data Collection Checklist

Table 4.1. Data collection checklist for the Turin case study (SPINA 3 – Corso Principe Oddone).

Dataset / File	Provider / source authority	What variables it contains	Why it matters for feasibility	How it is used in the workflow	Notes / limits
SPINA 3 CORSO ODDONE.docx	Project official documentation (client / project owner)	Site address; regulatory constraints; site area; buildable area; functional requirements and mix constraints.	Defines hard constraints and program boundary conditions that govern option feasibility.	Used as the primary constraint reference for all model prompts; cross-checked in Forma massing assumptions.	Scope is limited to what is documented in the project file.
Construction cost and time spend in Italy.xlsx	DEI srl Tipografia del Genio Civile (benchmark compendium, Northern Italy)	Indicative construction costs by building typology; indicative time schedules; example Gantt timelines.	Anchors early cost and schedule assumptions, reducing arbitrary "rule of thumb" inputs.	Used to populate baseline cost/time parameters and to sanity-check schedule narratives in reports.	Benchmark-based; not a project-specific bill of quantities.
Employment rate and Education level of Torino.xlsx	ISTAT official statistical datasets	Employment rate (by age groups, 2021–2024); education levels (age 9+, gender, 2021–2023).	Supports demand plausibility and target-segment definition, especially for residential and office uses.	Used as demand and socio-economic context inputs for market analysis narratives.	Aggregated indicators; not micro-level household panel data.
Gross Domestic Product.xlsx	ISTAT official statistical datasets	Italy GDP (2015–2024); Italy GDP per capita (2014–2023); Turin GDP per capita (2014–2022).	Provides macro context for affordability and investment climate assumptions.	Used to contextualize market cycle and scenario framing in the decision report.	Macroeconomic indicators cannot explain micro-location effects alone.
Households Size & Disposable income of the households.xlsx	ISTAT official statistics (household structure) + local income/tax statistics (postal-code level)	Household-size distribution (census zones); household disposable income trends (Piemonte, 2014–2023); Turin income/tax statistics by postal code (2023).	Informs target unit mix, affordability bands, and feasible pricing assumptions.	Used to justify residential unit mix logic and affordability-sensitive scenarios.	Some sub-series are spatially aggregated; cross-zone comparability must be stated.
OMI Price and Rent of D9 C8	Agenzia delle Entrate	Price and rent ranges for	Primary market evidence for	Used to derive baseline	OMI provides ranges;

Dataset / File	Provider / source authority	What variables it contains	Why it matters for feasibility	How it is used in the workflow	Notes / limits
C9 C10 C17 C19 2016-2025.xlsx	– OMI (official market observatory)	residential, office, and commercial uses by OMI zone (2016–2025).	price/rent assumptions and comparables in the target area.	price/rent ranges and to compare scenarios across neighboring zones.	transaction-level dispersion is not visible.
Piemonte and Torino Residential Price and Index 2016-2024.pdf	Agenzia delle Entrate – territorial directorates (regional residential market report)	Residential market indicators including transaction volumes (NTN) and market intensity (IMI), and related indices (2016–2024).	Supports cycle interpretation, absorption plausibility, and scenario sensitivity to market depth.	Used to support demand/absorption discussion and to contextualize price/rent evidence.	Document series; publication dates and reference periods must be tracked.
Populations of Torino within the city.xlsx	ISTAT official demographic datasets	Population and foreign-resident counts by sex and by spatial units (city / census areas; multiple years).	Shapes demand estimation, household formation assumptions, and service-load plausibility.	Used to define demographic baseline and to triangulate target segments for program choices.	Demographic indicators do not directly predict willingness-to-pay.
Surrounding area_data_2021 & Torino Municipal projections.xlsx	ISTAT census-zone indicators + municipal projections (as provided in dataset)	Census-zone profiles (education, employment, nationality, households, housing stock and vacancy) and official population projections for Turin (2030–2050).	Supports medium/long-term scenario framing and risk discussion (e.g., demand shifts, vacancy risk).	Used to develop scenario narratives and to justify sensitivity directions in the decision report.	Projection uncertainty must be explicitly noted in risk discussion.
Environmental & infrastructure visual data.pdf	Geoportale (official GIS portal) – compiled spatial layers	Maps/layers: flood hazard zones, access conditions, noise zoning, heritage constraints, postal-code context, and related spatial overlays.	Enables location-specific constraint interpretation and communicates spatial risks clearly.	Used as visual evidence for constraints; supports Forma checks and report figures/appendices.	Visual layers require careful referencing (layer name, date, scale).

The checklist is intentionally multi-scalar: it combines project-specific constraints with public statistics and market observatory outputs to reduce dependence on unverified assumptions. Where a variable is not available in the datasets above, the methodology requires that the report explicitly marks it as unavailable rather than imputing precise values without evidence (NIST, 2023; IVSC, 2024).

4.2.2 Principles for Data Organization and Upload

To keep LLM outputs auditable, data organization follows operational principles comparable to information-management discipline in built-environment delivery: clear naming, unambiguous versions, and a controlled common information environment logic. ISO 19650 emphasizes that information value depends on agreed naming, responsibilities, and approval workflows, which are directly relevant when an LLM is used as an evidence-processing layer rather than a mere writing tool (ISO, 2018a, 2018b).

Naming conventions and version clarity. File names are kept clean and descriptive, including dataset type, geography, time span, and version when applicable. Draft and final versions are not mixed. This reduces accidental use of outdated inputs and supports traceability.

Workbook structuring and internal organization. Spreadsheets are structured with clearly separated worksheets (e.g., raw data, processed indicators, and charts), and consistent units and definitions. This mitigates common spreadsheet risks such as silent overrides and hidden formula errors (Panko, 1998).

Language consistency. Each file is kept internally consistent in language (Italian or English) wherever possible. Mixed-language headers increase extraction errors and complicate later auditing, especially for regulatory terms and unit definitions.

File compression and format discipline. PDFs are compressed while preserving legibility, and all file types are kept as light as possible (targeting approximately ≤20 MB per file where feasible, as stated in the project synopsis). This reduces upload failures and incomplete reading.

Mitigating file-count limits. When platforms impose limits on the number of uploads, datasets are consolidated into fewer, well-structured files rather than arbitrarily dropping evidence. Consolidation actions are documented in the assumption log to preserve provenance.

Upload protocol and metadata. A short metadata header is maintained for each dataset (source authority, publication date, spatial coverage, and key variables). When sensitive information is involved, access and sharing follow GDPR-aligned principles of purpose limitation and minimization (European Union, 2016).

These principles are not primarily technical; they are governance mechanisms that help ensure that any model output can be traced back to the evidence it claims to use. This is consistent with valuation and feasibility standards, which treat documentation, transparency, and reviewability as fundamental requirements (IVSC, 2024; RICS, 2024).

4.3 Testing and Selection of Large Language Models

Given the documented variability of LLM outputs, model choice was treated as an empirical selection problem rather than a preference. Stress testing was therefore used to compare leading models under realistic constraints relevant to this project: multi-file evidence ingestion, hard-constraint compliance, and the production of an auditable decision report.

4.3.1 Large Language Model Testing Methodology

The stress test was designed as a controlled, single-run task: each model received the same dataset bundle and the same prompt, and was asked to generate one end-to-end real-estate decision report in Word format using Deep Research mode. In this thesis, the term "decision package" is operationalized as the completeness of required analytical components within that single report artefact (market analysis, financial analysis, risk analysis, and other essential components). The report output was then evaluated using a structured rubric (Section 4.3.3).

To ensure comparability, the prompt specified a consistent internal structure for the decision report. Beyond standard market/finance/risk sections, the report was required to articulate project-relevant design and feasibility choices, including: (i) a recommended residential unit mix (e.g., studio, 1-bedroom, 2-bedroom, 3-bedroom); (ii) a recommended commercial typology and operational model; (iii) an integration strategy for residential, office, and commercial components (e.g., podium-and-tower, residential above commercial/office, standalone office buildings); (iv) a site layout and massing strategy explicitly linked to key evaluation criteria (noise exposure, commercial attractiveness, proximity to services, views/landscape, and flood-risk mitigation); (v) service infrastructure and technical provisions (parking, MEP, utilities); and (vi) an indicative construction timeline supported by a detailed Gantt chart. Where models could not produce figures directly, they were expected to provide clear placeholders for charts, diagrams, and conceptual sketches.

To keep the evaluation aligned with architecture and feasibility practice, the rubric dimensions were defined in plain professional terms: (1) perception and spatial reasoning (how well spatial constraints and site logic are understood), (2) evidence and data reliability (whether claims are grounded in provided data and clearly bounded), (3) decision quality (clarity, coherence, and actionable recommendations under constraints), (4) coverage and insight (whether the report covers essential components without superficiality), and (5) usability and implementation constraints (workflow friction, practical limits, and output usability in a real project setting) (NIST, 2023; IVSC, 2024; RICS, 2024).

Limitations. First, reproducibility is constrained by platform updates and model drift: the same prompt may yield different outputs as models change over time. Second, access tiers and subscription levels affect available features (e.g., project folders, file limits, and mode availability), meaning the results reflect deployability under the tested access conditions rather than an abstract benchmark. Third, multimodality differs across platforms; models that cannot reliably process images or complex PDFs are disadvantaged for built-environment evidence. Finally, the rubric improves transparency but does not eliminate expert judgement, particularly in interpreting hard-constraint violations and auditability gaps (NIST, 2023; ISO, 2018a, 2018b).

4.3.2 Testing Procedures for Major Large Language Models

All proprietary models were tested under paid or provisioned access where required, while open or free models were tested under their publicly available conditions. The procedures were kept consistent: upload the same evidence package, run the same Deep Research prompt, and record platform constraints and observed failure modes in a neutral manner.

Where platforms supported a project workspace, the environment was configured with a standardized project instruction defining the role (academic feasibility/report writing), the evidence discipline (do not invent facts; state when data are unavailable), and the required output structure. This reduced contamination from unrelated chat context and improved comparability across runs.

Table 4.2. Practical constraints observed during LLM stress testing (platform features and limits).

Model	**Access condition (as tested)**	**Project workspace available (as tested)**	**Modes used / available (as tested)**	**Upload limits observed (count/size)**	**Multimodal handling observed**	**Key constraints / failure modes (observed)**
ChatGPT 5.2	Paid Plus plan (author's test setup)	Yes (project workspace used)	Deep Research for report generation; extended reasoning for structured extraction and modelling	Observed cap: ≤25 files in project workspace	Handled mixed DOCX/XLSX/PDF inputs in this study	Most complete and stable reports; remaining issues include occasional idealized assumptions and visual placeholders requiring human verification.
Gemini3.0	Student access plan (author's test setup)	No project workspace in the tested environment	Deep Research used for report generation; limited mode conversion in the test workflow	Observed practical constraint: file-count limits required selective upload	Handled mixed DOCX/XLSX/PDF inputs in this study	Hard-constraint violation in the first run (ignored project height cap); outputs remained relatively short/shallow compared with the best-performing model.
Claude Opus 4.5	Paid Pro plan (author's test setup)	Yes (project workspace used), plus chat-upload workaround	Deep Research used; iterative refinement constrained by platform limits	Upload constraints required consolidation, with a max file size of 31MB each	Handled mixed DOCX/XLSX/PDF inputs in this study	Very high variance between runs; strict conversation-turn limits reduced practical deployability for iterative feasibility work.
DeepSeek	Free/open access (author's test setup)	Not available / not used	No Deep Research mode in the test workflow	Upload up to 50 files simultaneously, with a max file size of 100 MB each	Limited for image-based/complex evidence in the test log	Could not satisfy stress-test requirements; low information-extraction rate was observed (13% reported in test notes).

Model	Access condition (as tested)	Project workspace available (as tested)	Modes used / available (as tested)	Upload limits observed (count/size)	Multimodal handling observed	Key constraints / failure modes (observed)
Grok4.1	Paid access (author's test setup)	Project workspace existed, but the report workflow could not be integrated cleanly	Deep Research-style workflow was not integrated with the project workspace in the test log	Not specified in the test log	Handled mixed DOCX/XLSX/PDF inputs in this study	Insufficient capacity and overly informal reporting style for audit-critical feasibility documentation in the test results.
Qwen3-Max	Free access (author's test setup)	Yes (project supported)	Deep Research used	Observed limits: max 5 files; total size cap 20 MB	Handled mixed DOCX/XLSX/PDF inputs in this study	Severe upload constraints required aggressive consolidation; outputs tended to be shorter/less detailed, increasing the risk of silent omissions.

ChatGPT 5.2 was tested using a project folder to isolate the case-study context from unrelated chat history, with the evidence package uploaded into the project workspace (with a noted cap of 25 files). Two report-generation approaches were attempted: a more granular approach in which the model first analyzed each dataset in extended reasoning mode and then generated the final decision report in Deep Research mode; and a more direct approach that prioritized producing the report after upload. Across both outputs, ChatGPT 5.2 produced the most complete and stable reports among the tested models, though it still relied on placeholders for visuals and occasionally introduced idealized elements that require human verification.

Gemini3.0 was constrained by the lack of a project workspace in the tested environment and by limited mode flexibility (Deep Research could only be run in a separate chat environment without conversion). The first report showed a major hard-constraint failure (residential towers up to approximately +15 floors despite a 5–7 floor cap noted in the project constraints) and provided relatively shallow financial and risk treatment. A second attempt improved technical detail and reduced obvious errors, but the output remained short and less traceable.

Claude Opus 4.5 supported a project workflow and could accept the evidence package via a project folder; however, the first output under-used the uploaded evidence and resembled a generic market survey rather than a project-specific feasibility report. A second attempt—requiring a different upload strategy—produced a substantially better feasibility report. In practical terms, the most significant constraint was the platform's strict conversation-turn limitation, which prevented iterative refinement after only a few rounds and therefore reduced deployability in a real workflow.

DeepSeek and Grok were tested but excluded from rubric scoring because they could not reliably satisfy the stress-test requirements under the available modes and constraints. DeepSeek lacked Deep Research mode and multimodal capacity, and its reported information extraction rate was too low for audit-critical feasibility work. Grok offered a Deep Research mode but separated it from the project workflow and showed insufficient capacity and overly informal reporting style for professional feasibility documentation.

Qwen3-Max offered a project feature but imposed strict upload constraints (maximum five files and a size cap of 20 MB). This required heavy file consolidation, which itself increases the risk of losing provenance and introducing silent omissions. Within these limits, Qwen produced outputs that were sometimes free of obvious errors but often lacked depth and detail.

4.3.3 Scoring Sheet and Scoring Results Discussion

All completed decision reports were scored using a rubric designed to make qualitative evaluation explicit and auditable. Each rubric item is scored on a 0–2 scale (0 = not met/clearly wrong; 1 = partially met; 2 = fully met/verifiable and consistent). Dimension scores are then normalized to a 0–10 scale using: Dimension score = (points earned / maximum points) × 10. The overall score is computed as the weighted sum of the five dimension scores. This approach is aligned with governance expectations that require transparent criteria, evidence notes, and traceable judgement (NIST, 2023; IVSC, 2024; RICS, 2024).

Rubric weights (Table 4.3) are taken directly from the scoring workbook (consistent across the per-report "Summary-*" sheets, e.g., "Summary-GPT01"), and overall model averages/rankings (Table 4.4) are taken from the workbook summary sheet "Summary-total".

Table 4.3. Rubric dimensions and weights used for scoring the decision reports.

Dimension	Weight	Interpretation in feasibility practice
1. Perception and Spatial Reasoning	0.20	Quality of spatial understanding, constraint interpretation, and built-environment logic.
2. Evidence and Data Reliability	0.25	Use of provided datasets, explicit uncertainty bounds, and avoidance of unsupported claims.
3. Decision Quality	0.25	Actionable recommendations, coherent logic, and scenario-sensitive reasoning under constraints.
4. Coverage and Insight	0.20	Completeness across market/finance/risk and depth beyond generic statements.
5. User Experience and Engineering Implementation	0.10	Practical workflow viability: limits, exports, integration friction, and usability.

Table 4.4. Overall ranking and model scores from the stress-test rubric (average across two reports per model).

Rank	Model	Overall model score (0–10)	Report IDs (2 runs)	Report weighted scores
1	ChatGPT 5.2	7.38	Chatgpt01 Spina 3 Oddone Redevelopment – Comprehensive Development Strategy; Chatgpt02__Spina 3 – Corso Oddone__ Real Estate Development Strategy Report	7.18; 7.58
2	Claude Opus 4.5	6.05	Claude 01 SPINA 3 Corso Oddone development comprehensive feasibility research; Claude 02 SPINA3_Comprehensive_Feasibility_Report	3.68; 8.42
3	Qwen3-Max	4.21	Qwen 01 A Strategic Real Estate Development Plan for an Urban Regeneration Project in Turin, Italy; Qwen 02 2, San Donato, Turin	3.35; 5.08
4	Gemini3.0	4.11	Gemini01 Real Estate Report_ Spina 3 Project; Gemini02 Real Estate Development Report Generation	3.90; 4.32

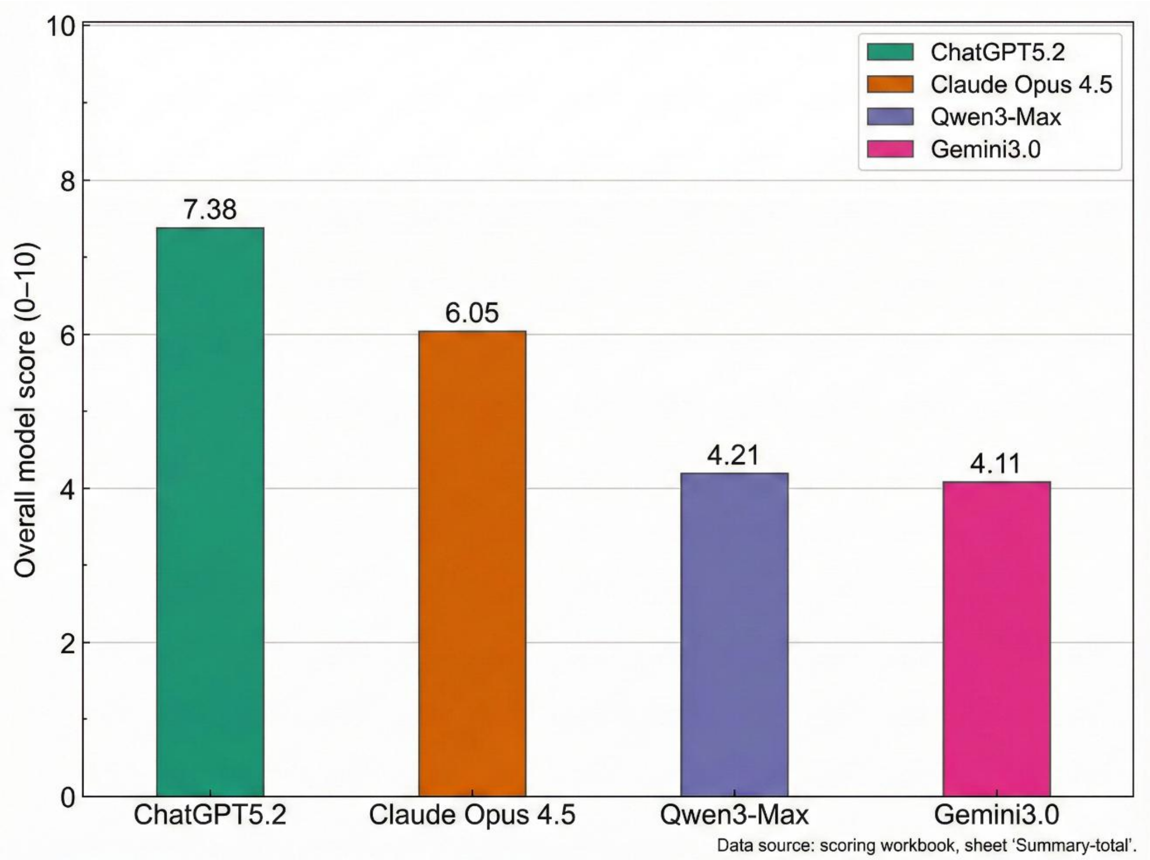

Figure 4.2. Overall model score comparison (weighted average across two decision-report runs).

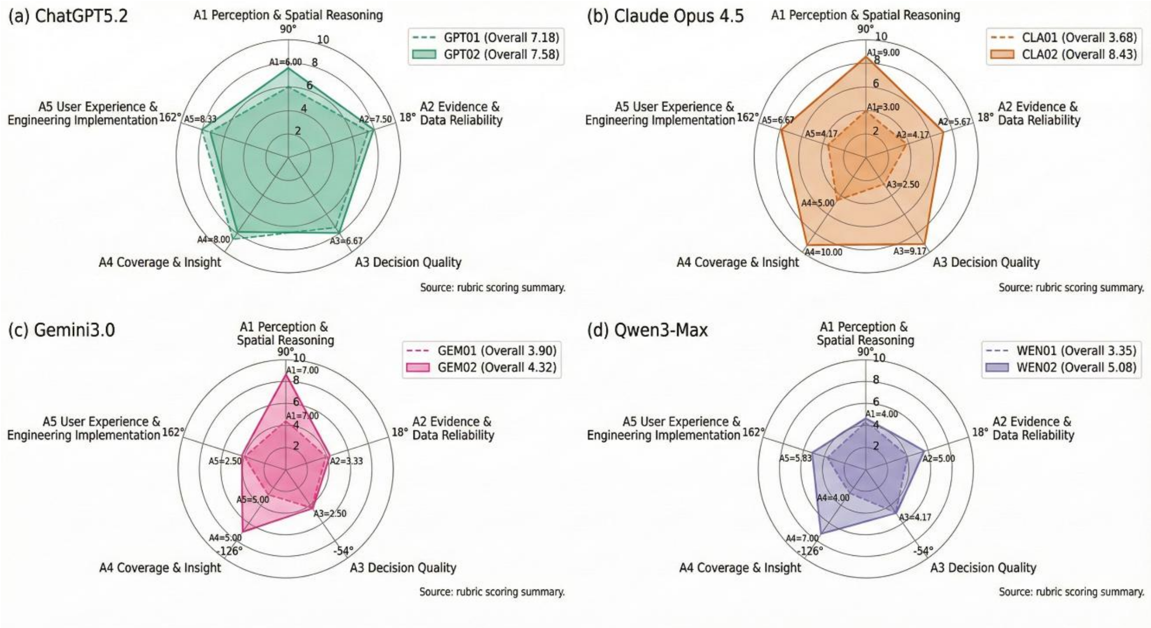

Figure 4.3. Dimension-profile comparison across models (average of two runs per model).

The scoring results indicate that ChatGPT 5.2 achieved the highest overall performance and the most consistent decision-report quality across two runs (scoring workbook sheet "Summary-total"; see Table 4.4). Claude Opus 4.5 ranked second on average, but with very high variance between its two reports, highlighting sensitivity to workflow conditions and platform constraints. Qwen3-Max and Gemini3.0 ranked third and fourth respectively in the rubric summary, reflecting weaker sourcing/traceability and/or practical constraints that reduced completeness under the tested conditions (Tables 4.2 and 4.4).

Based on the rubric results and the prior SWOT findings discussed in Chapters 2–3, this thesis adopts ChatGPT 5.2 as the primary LLM for evidence synthesis, constraint checking narratives, assumption-log extraction, and the production of the core feasibility report. For visualization support—specifically fixing or replacing charts and figure placeholders—the method pairs the Gemini ecosystem with Nano Banana Pro as the primary image-generation tool, because the stress testing and workflow experience indicated that visual deliverables were a recurring weakness in text-only outputs and that a dedicated multimodal tool improves report readability without changing the underlying governed numbers.

4.4 Final Methodology for Integration of AI into the Real-Estate Decision-Making, Based on Test Scores and SWOT Analysis

The final integrated methodology adopts a deliberately constrained toolchain to balance capability with deployability under EU/Italy project conditions. The tool chain is executed in the following sequence, with explicit output artefacts and version control at each step:

I. ChatGPT 5.2 (Deep Research + reasoning): Data collection synthesis; compliance framing; market analysis; concept narrative; extraction of a single-source assumption set (SSOT/SSOP) into an assumption log; financial modelling (returns and cash-flow logic) and sensitivity/risk analysis. Outputs: named and versioned Excel/PDF exports of assumptions and model results, plus a draft decision report.

II. Autodesk Forma: Technical feasibility checks, massing and modelling, and early environmental checks. Outputs: exported snapshots, key metrics, and documented assumptions used for geometry and rule checks.

III. ChatGPT 5.2 (synthesis and governance): Update the assumption log based on Forma outputs; regenerate and reconcile model outputs; synthesize the feasibility report narrative with explicit traceability notes. Outputs: revised, versioned Excel/PDF artefacts and a consolidated Word feasibility report draft.

IV. Nano Banana Pro (Gemini ecosystem): Replace or refine chart and figure placeholders; convert key tables to clear visuals (without changing the numeric SSOT). Outputs: image files inserted into the report and archived alongside the versioned spreadsheets.

V. Human final review and delivery: Final verification of hard-constraint compliance, reasonableness checks on market/cost assumptions, and editorial quality control. Deliverables are packaged with clear version identifiers and a minimal audit trail (assumption log, inputs list, and output register) (ISO, 2018a, 2018b; IVSC, 2024; RICS, 2024).

Version control is treated as a mandatory governance layer rather than an optional best practice. Each major run produces a named bundle of artefacts (assumption log, model outputs, and report draft) with a version identifier that is recorded in a short output register. This supports later review, comparison across concept options, and transparent reporting of uncertainty and changes—requirements that align with both built-environment information management and professional valuation documentation expectations (ISO, 2018a, 2018b; IVSC, 2024; RICS, 2024; European Union, 2024).

5. Empirical Application to a Real Case Study

This chapter applies the AI-integration methodology established in Chapter 4 to a real redevelopment case in Turin. Instead of treating AI as a standalone "generator," it operationalizes AI as part of an auditable, version-controlled workflow for concept-stage feasibility. The process begins by consolidating heterogeneous evidence into a curated dataset and stabilizing constraints and assumptions through a single source of truth (SSOT). On this basis, a single decision report is generated, checked, and revised through explicit human audit points, so that the reasoning chain remains traceable. The chapter then extends the workflow downstream: the decision report functions as a translation layer from feasibility logic to spatial actions, allowing the three concept options to be tested through measurable massing, phased planning, and early environmental checks using Autodesk Forma as the geometry-and-quantity SSOT. The objective is not to claim final design resolution or regulatory approval, but to demonstrate a defensible bridge from constraints and spatial logic to comparable, carry-forward assumptions for subsequent financial modelling and comparative decision-making.

5.1 Case Study Introduction

SPINA 3 – Corso Principe Oddone is located in the northern inner ring of Turin, within the San Donato district, and forms part of the wider urban transformation corridor known as the Spina Centrale. The case-study site is a former railway-related area associated with the Passante Ferroviario and the subsequent reconfiguration of surface infrastructure, where the historic rail trench has been progressively covered and reconnected to the surrounding urban fabric through the new Viale della Spina system (Comune di Torino, 1995; Comune di Torino, 2000; MuseoTorino, n.d.).

At the neighborhood scale, the site sits between major movement corridors (Corso Principe Oddone and the broader Spina axis) and the Dora River landscape system. It is proximate to Parco Dora—an emblematic post-industrial regeneration project within Spina 3—and to Environment Park, a research and innovation cluster that influences the area's identity and demand profile (MuseoTorino, n.d.; SPINA 3 CORSO ODDONE, n.d.). The perimeter and immediate context referenced in this thesis follow the project brief: Via Savigliano, Corso Principe Oddone, Corso Vigevano, Corso Umbria, and Piazza Baldissera (SPINA 3 CORSO ODDONE, n.d.).

Figure 5.1 Indicative satellite view and boundary of the "Spina 3 – Oddone" site along Corso Principe Oddone (Turin). Source: Italian Trade Agency (ICE), 2016.

From a planning and governance perspective, the area is designated as an Urban Transformation Zone (Zona di Trasformazione Urbana, ZUT) 4.13/2 "Spina 3 – Oddone", which requires implementation through an executive planning instrument (SUE) and a related convention, consistent with Turin's PRG framework for transformation areas (FS Sistemi Urbani, 2025). In the SSOT concept report used for this thesis, the total site area (Superficie Territoriale, ST) is 109,183 m², and the feasible development capacity is framed at approximately 46,006 m² of gross floor area (GFA) under the stated rules and constraints (Feasibility Study – Constraints, Market & Concepts, 2025).

Key quantitative constraints and obligations reported in the project development brief include: (i) a mandated mixed-use program with a minimum residential share (≥40% of GFA) and capped non-residential components (ASPI≤20％ of GFA and Eurotorino≤40％ of GFA); (ii) substantial land allocation to public services and open space, including approximately 56,649 m² of areas to be dedicated to standards and public realm; and (iii) required extensions of the Via Dronero and Via Ceva street corridors through the site to ensure permeability and integration with the existing grid (Feasibility Study – Constraints, Market & Concepts, 2025; FS Sistemi Urbani, 2025). The SSOT further notes an existing protected/heritage building to be preserved and integrated, which conditions both massing strategies and phasing choices (Feasibility Study – Constraints, Market & Concepts, 2025).

In addition, the project development brief records envelope-related constraints, including maximum building heights expressed in stories (e.g., 5 stories for interior zones and up to 7 stories along primary edges, depending on sub-area), as well as statutory parking ratios to be verified during later design development (Feasibility Study – Constraints, Market & Concepts, 2025).

Figure 5.2 Bird's-eye view context of Corso Principe Oddone and Piazza Baldissera, illustrating the infrastructural character and urban edges. Source: Città di Torino, 2023.

These constraints make the site suitable for empirically testing the Chapter 4 methodology: it requires combining heterogeneous evidence (planning rules, market context, and spatial feasibility) into an auditable decision pathway, while remaining at concept-stage where design alternatives can still be tested through rapid iteration. The following sections therefore apply the two-chat workflow and the selected toolchain to generate a versioned SSOT report, and then validate three proposals against the hard constraints through parametric modelling and early environmental checks.

At the metropolitan scale, the case study also sits within Turin's broader policy emphasis on urban regeneration, public-space investment, and climate-oriented upgrades supported through national and EU funding streams such as the PNRR. While the thesis does not attribute specific PNRR projects to the Spina 3 – Oddone site, the program context provides a relevant background for the emphasis on mixed-use transformation, public realm, and environmental performance checks at concept-stage (Comune di Torino, 2025).

5.2 Workflow Logic and Principle

The empirical workflow was implemented as a two-chat structure aligned with the methodology in Chapter 4: an input-side session for specifying the prompt logic and uploading the curated dataset set, and an output-side session for generating and iterating the report artefacts. This separation is operational rather than conceptual: it responds to platform constraints (context-window limits and compute budgeting) and reduces "context noise", while improving traceability by keeping the prompt history and outputs clearly separable and versioned (ISO, 2018a; ISO, 2018b). Within this chapter, the SSOT (Single Source of Truth) refers to the single, versioned document that consolidates the constraints, assumptions, and concept definitions used downstream.

Within this structure, the human researcher functions as reviewer and auditor. Each intermediate artefact was checked for internal consistency, explicit sourcing, and compliance with the hard constraints before it was accepted as an input to the next step. This audit checkpoint reflects the governance logic discussed in Chapter 4: credibility is established through verification and transparent documentation, not through model confidence alone (Chapter 4).

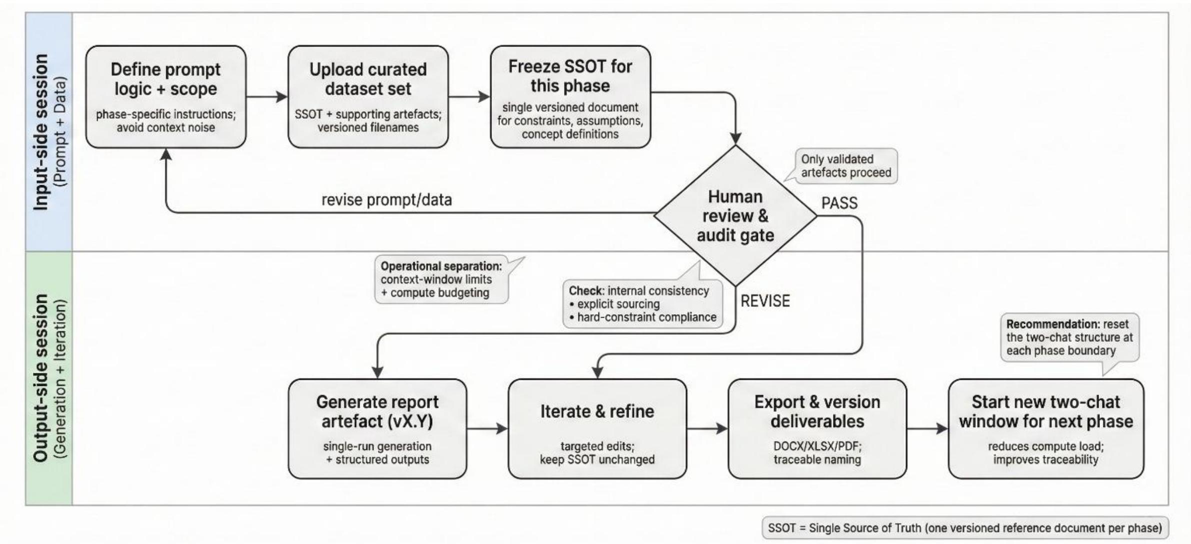

Figure 5.3 Two-chat workflow for the empirical case study, showing an input-side prompt/data session and an output-side generation session, with a human audit checkpoint that validates each versioned artefact before proceeding. Source: Author's diagram, 2025.

5.3 Workflow of Methodology for AI Integration Across the Entire Real Estate Decision-Making Process

Section 5.2 documents how the integrated toolchain selected in Chapter 4 was executed in practice. The emphasis is on traceability: each step produces identifiable, versioned artefacts (report drafts, assumption logs, and model exports), and the human reviewer acts as an auditor who checks constraints and evidence use before downstream modelling. The workflow therefore relies on controlled inputs and outputs rather than on conversational memory.

5.3.1 Concept Report

The first empirical step operationalized the "one-shot" Deep Research output described in Chapter 4: using a single fixed prompt and the curated dataset set, the platform generated a standalone Word report that functions as the project's Single Source of Truth (SSOT) for the subsequent modelling gates. In this case study, the resulting artefact is the document "1. Cover & Version.docx" (titled *Feasibility Study – Constraints, Market & Concepts*), which consolidates:

hard planning and administrative constraints; a targeted market scan for Turin/Spina 3; three comparable concept options for redevelopment (Feasibility Study – Constraints, Market & Concepts, 2025).

Crucially, the output is not a structured "decision package" generated through multiple linked files, but rather a single, end-to-end generated narrative report, excluding financial, sensitivity, and risk analyses. For auditability, the report was reviewed before proceeding to parametric modelling: the author checked that key constraints (e.g., zoning designation, land-use mix rules, and site-area and GFA values) were explicitly stated, internally consistent, and traceable to the cited sources, and that any missing data were marked as not specified rather than implicitly assumed (Chapter 4; Feasibility Study – Constraints, Market & Concepts, 2025).

The SSOT report proposes three concept archetypes that translate the same regulatory envelope into distinct spatial and program logics:

Concept 1 prioritizes housing and green public realm;

Concept 2 intensifies the tertiary component to form a mixed-use sub-center;

Concept 3 introduces an innovation-campus anchor combined with residential delivery.

Table 5.1 SSOT summary of Concept 1–3 (program intent, key indicators, and defining features). Source: Feasibility Study – Constraints, Market & Concepts (2025).

Concept	**Intended program mix (as stated)**	**Key quantitative indicators (as stated in SSOT)**	**Defining spatial / strategic features (SSOT summary)**
Concept 1: Residential-Led Green Quarter	Use Mix: ~60% Residential; ~15% Office; ~10% Retail; ~15% Civic / community services.	≈46,000 m² total GFA; residential GFA ≈27,600 m²; ≈300 dwelling units (rounded in SSOT).	Perimeter-block massing with internal green spine/park; local-services focus; modest tertiary component.
Concept 2: Balanced Urban Hub	Use Mix: ~40% Residential; 40% Tertiary (Office/R&D); 20% Commercial / services.	≈46,000 m² total GFA; offices ≈18,500 m²; retail ≈9,000 m²; ≈200 dwelling units (approx.).	Higher-intensity mixed-use center with central plaza and stronger employment component; phasing depends on office absorption/anchors.
Concept 3: Innovation Campus & Housing	Use Mix: ~50% Residential; ~30% Tertiary; ~10% Commercial; ~10% specialty/education component (as described).	Residential GFA ≈23,000 m²; ≈250 dwelling units (approx.); student residence ≈100 beds (as described).	Innovation-campus cluster near Environment Park; porous public realm and promenade; relies on institutional/partner commitment.

The three concepts were then operationalized as three Autodesk Forma proposals, calibrated to match the SSOT area proportions and spatial logics, in order to verify technical feasibility at concept-stage (Section 5.3.2).

5.3.2 Modeling and Technical Feasibility Study

In the second empirical step, Autodesk Forma was used to translate the three concept proposals from the Concept Report into concept-stage massing models (Proposals 1–3). The purpose was to validate, in geometry, what cannot be safely validated in text alone: that each option can satisfy the required program proportions and area constraints while remaining a plausible built form. Each proposal was modelled to reflect the program distributions and spatial logic defined in the Concept Report, and then evaluated using Forma's Area metrics and early environmental analyses. This step operationalizes the division of labor set out in Chapter 4: the LLM supports evidence synthesis and option articulation, while a parametric platform supports measurable spatial validation (Autodesk, n.d.-b). In this chapter, each Concept was implemented as a corresponding Proposal in Forma (Proposal 1 = Concept 1; Proposal 2 = Concept 2; Proposal 3 = Concept 3).

Forma's Area metrics analysis provides numerical feedback required for concept-stage feasibility screening, including site area, building coverage, gross floor area, internal areas, and program splits. In this thesis, these metrics are used as a verification layer against the SSOT constraints in the Concept Report and as a structured basis for the SSOT assumption log that will feed subsequent financial modelling. Importantly, the metrics confirm whether an option is 'within bounds' at the massing level; they do not substitute for detailed measurement protocols required at later design gates (Autodesk, n.d.-a; RICS, 2015; RICS, 2023).

Table 5.2 Autodesk Forma area metrics for Proposals 1–3 (exported from Area metrics; values in m2 unless noted). Source: Author’s Autodesk Forma export, 2025 (Area metrics).

Proposal 1

Figure 5.4 Autodesk Forma massing models for Proposals 1, shown at consistent scale and viewpoint for comparative review.

Metric	Total	Residential	Commercial	Office	Public
Site area (m2)	46734				
BC (m2)	8734				
GFA (m2)	46006	29454	4717	7075	4761
GIA (m2)	41405	26508	4245	6367	4285
NIA (m2)	37265	23858	3820	5731	3856
BFA (m2)	11557				
GSA (m2)	20766				
Unit count	308				

Proposal 2

Figure 5.5 Autodesk Forma massing models for Proposals 2, shown at consistent scale and viewpoint for comparative review.

Metric	Total	Residential	Commercial	Office	Public
Site area (m2)	46734				
BC (m2)	9095				
GFA (m2)	46006	18948	9170	17888	0
GIA (m2)	41405	17053	8253	16099	0
NIA (m2)	37265	15348	7428	14489	0
BFA (m2)	12004				
GSA (m2)	20783				
Unit count	201				

Proposal 3

Figure 5.6 Autodesk Forma massing models for Proposals 3, shown at consistent scale and viewpoint for comparative review.

Metric	Total	Residential	Commercial	Office	Public
Site area (m2)	46734				
BC (m2)	9180				
GFA (m2)	46006	23012	4615	13812	4566
GIA (m2)	41406	20711	4154	12431	4110
NIA (m2)	37265	18640	3738	11188	3699
BFA (m2)	12004				
GSA (m2)	20783				
Unit count	255				

Forma reports area metrics using project- and region-specific settings, and it allows users to adjust calculation rules and ratios for key figures (Autodesk, n.d.-a). For this reason, the metrics below are defined in relation to (i) their role within the Forma export and (ii) widely used built-environment measurement conventions:

- Site area: the site boundary area used in the Forma model for the proposal. (Autodesk, n.d.-a).
- BC (Building coverage): the ground-level building coverage area reported by Forma (used to assess site coverage/footprint intensity). (Autodesk, n.d.-a; RICS, 2023).
- GFA (Gross floor area): the gross floor area reported by Forma overall and by assigned function (program split). (Autodesk, n.d.-a; RICS, 2023).

- GIA (Gross internal area): the internal area measured to the internal face of the perimeter walls at each floor level (a common convention for gross internal area). (Autodesk, n.d.-a; RICS, 2015).
- NIA (Net internal area): the usable internal area within a building measured to the internal finish of perimeter walls, excluding non-usable common areas according to the adopted basis. (Autodesk, n.d.-a; RICS, 2015).
- BFA (Building Footprint Area) refers to the horizontal footprint of the above-ground building mass on the site, reported by Autodesk Forma as part of its Area metrics export. At concept-stage, BFA is useful for checking coverage logic against open-space and permeability objectives, but it remains sensitive to modelling assumptions and massing granularity (Autodesk, n.d.-a).
- GSA (Green Space Area) refers to the portion of the site explicitly modelled as green/open landscape surface (e.g., planted areas) in the Autodesk Forma scheme. In this thesis it is used as an indicative proxy for early open-space performance; however, it does not replace detailed landscape design or statutory 'urban standards' calculations at later gates (Autodesk, n.d.-a).
- Unit count: the total number of residential units reported for the proposal as part of the proposal summary. (Autodesk, n.d.-b).

Across the three proposals, the total GFA is consistent at 46,006 m2, aligning with the maximum buildable GFA stated in the Concept Report. Differences appear primarily in coverage and unit counts, reflecting the differing spatial logics of the three concepts (Feasibility Study – Constraints, Market & Concepts, 2025).

After confirming the area metrics and program splits, Forma's environmental analyses were used to support early-stage technical feasibility screening. These analyses are intended for concept-stage decision support: they enable consistent comparisons between options under a shared modelling set-up, but they remain assumption-driven and sensitive to the choice of dates, climate inputs, default material parameters, and geometric resolution. Accordingly, the outputs are interpreted as comparative signals (e.g., identifying relative risks and opportunities) rather than as evidence of regulatory compliance or detailed performance (Autodesk, n.d.-b).

Sun hours: indicates hours of direct sunlight on selected dates, supporting early evaluation of outdoor-space quality and façade exposure. Indicative only; results depend on the selected dates/time interval and simplified geometry. (Autodesk, n.d.-c; Autodesk, n.d.-b).

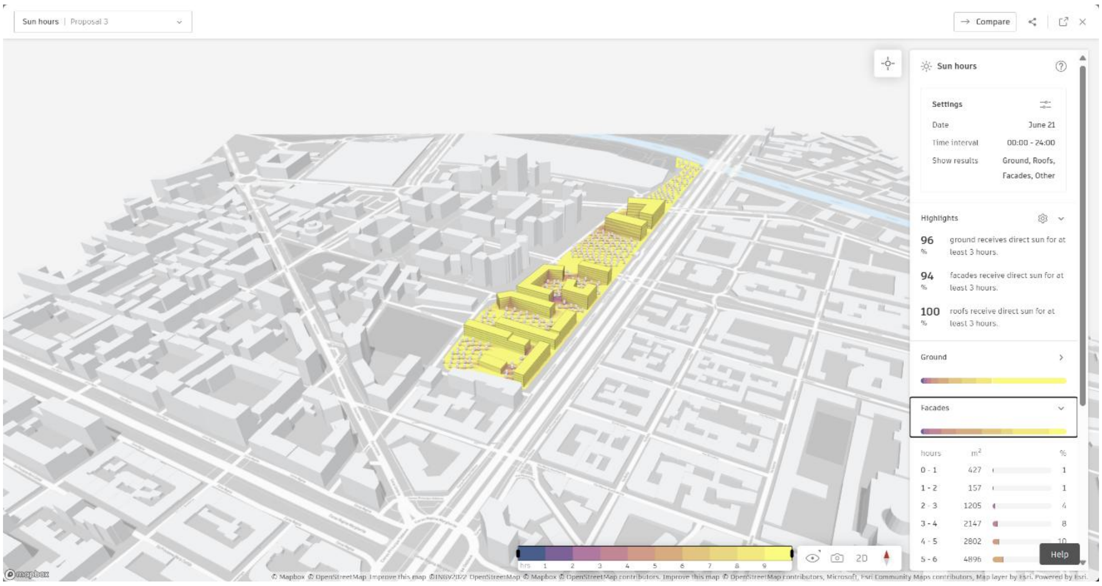

Figure 5.7 Example Autodesk Forma environmental-analysis dashboard outputs for Proposal 3, illustrating concept-stage comparative checks (Sun hours). Source: Author's Autodesk Forma outputs, 2025.

Daylight potential: provides an early indicator of daylight access on façades, supporting massing decisions that reduce self-shading. A potential-based assessment; it does not replace detailed daylight simulations at later gates. (Autodesk, n.d.-b).

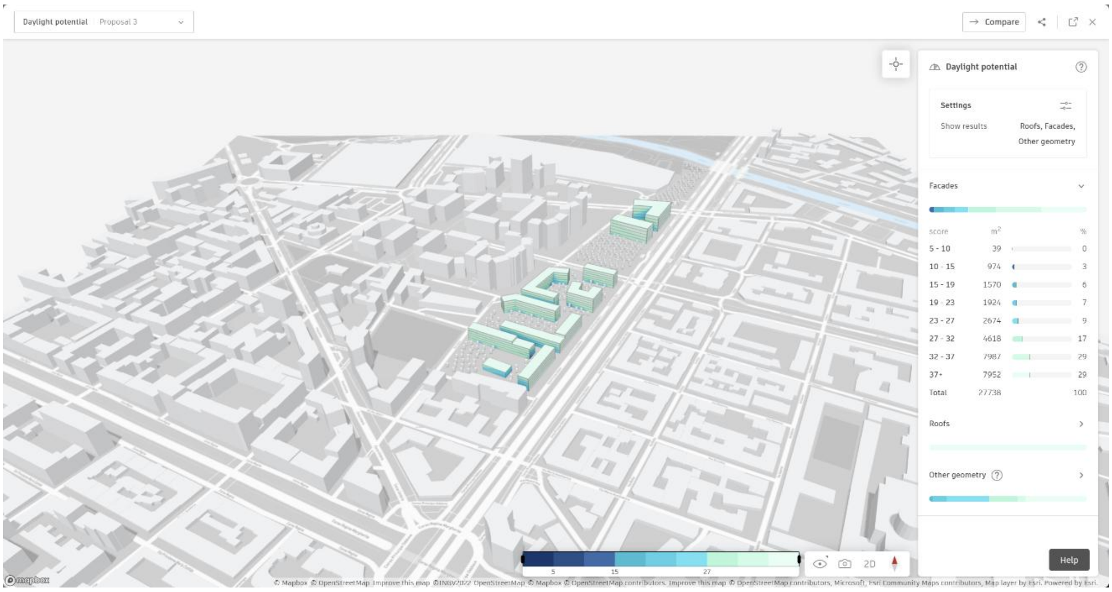

Figure 5.8 Example Autodesk Forma environmental-analysis dashboard outputs for Proposal 3, illustrating concept-stage comparative checks (Daylight potential). Source: Author's Autodesk Forma outputs, 2025.

Embodied carbon: offers a high-level comparison signal for embodied carbon implications across alternatives during concept design. Assumption-driven; sensitive to default material and structural assumptions. (Autodesk, n.d.-d).

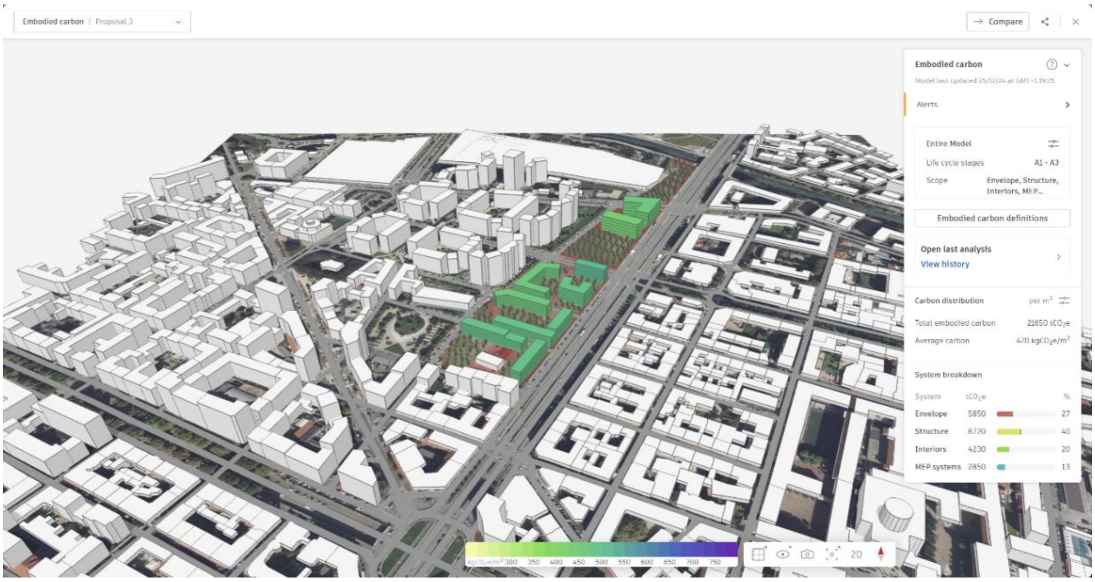

Figure 5.9 Example Autodesk Forma environmental-analysis dashboard outputs for Proposal 3, illustrating concept-stage comparative checks (Embodied carbon). Source: Author's Autodesk Forma outputs, 2025.

Microclimate: provides an indicative view of outdoor comfort conditions based on sun/wind interactions and location data. Should be complemented by detailed studies once the design is fixed. (Autodesk, n.d.-d; Autodesk, n.d.-b).

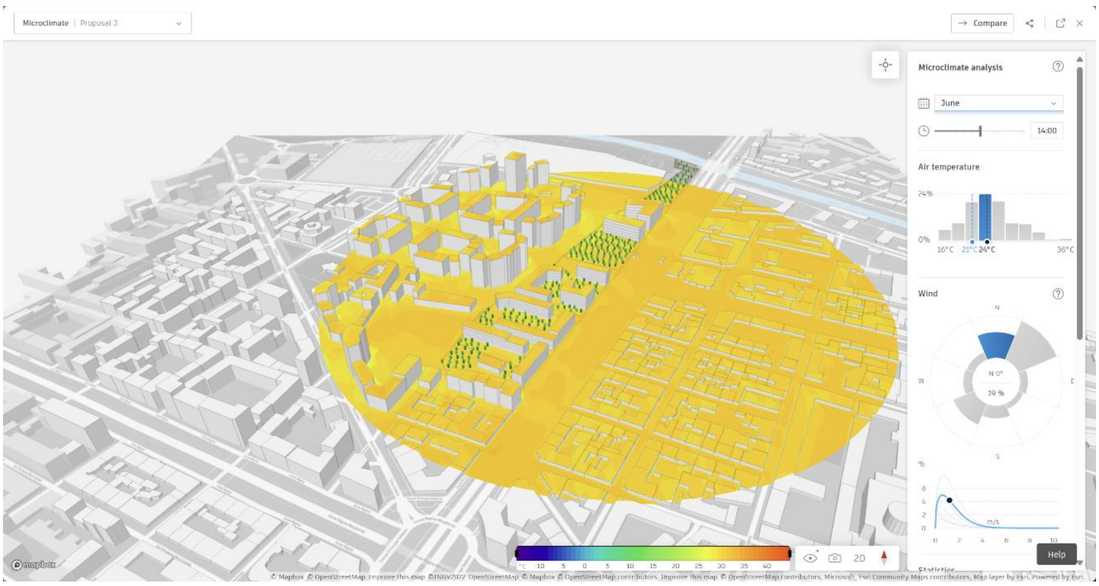

Figure 5.10 Example Autodesk Forma environmental-analysis dashboard outputs for Proposal 3, illustrating concept-stage comparative checks (Microclimate). Source: Author's Autodesk Forma outputs, 2025.

Solar energy: estimates solar energy potential and supports early evaluation of photovoltaic opportunity and surface suitability. Sensitive to surface modelling and assumed PV placement. (Autodesk, n.d.-b).

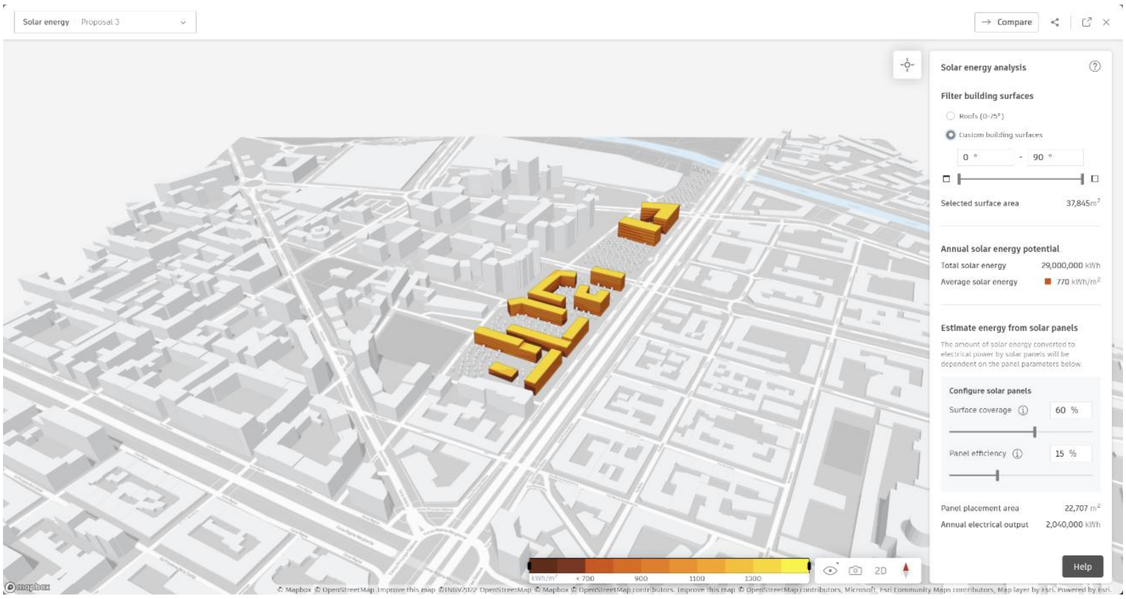

Figure 5.11 Example Autodesk Forma environmental-analysis dashboard outputs for Proposal 3, illustrating concept-stage comparative checks (Solar energy). Source: Author's Autodesk Forma outputs, 2025.

Wind: supports early identification of wind comfort risks around massing, including rapid pre-analysis modes. Requires validation if used for documentation or compliance. (Autodesk, n.d.-b).

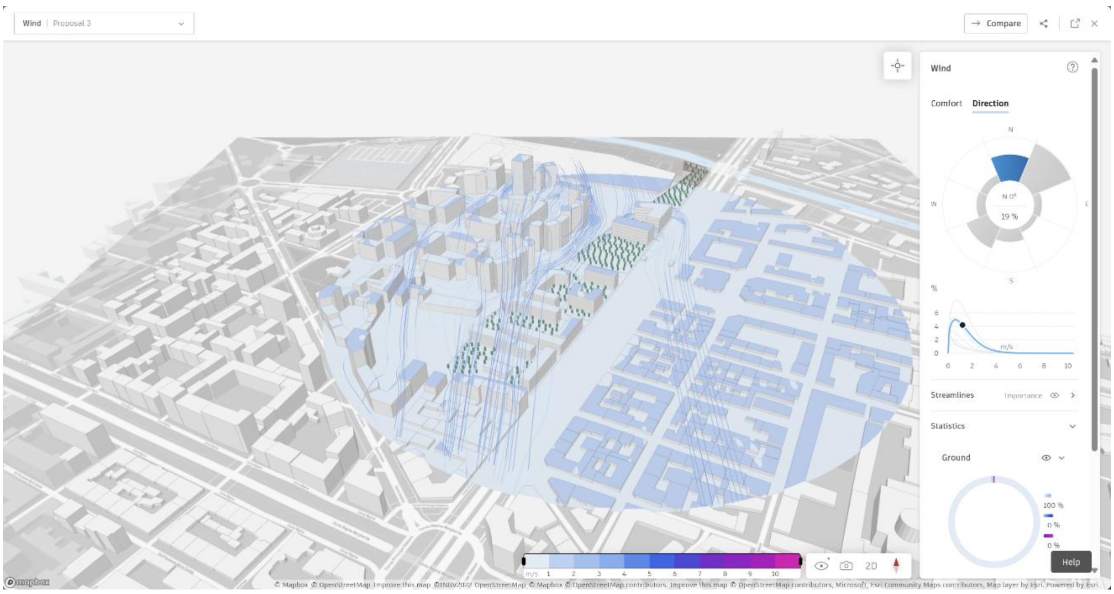

Figure 5.12 Example Autodesk Forma environmental-analysis dashboard outputs for Proposal 3, illustrating concept-stage comparative checks (Wind). Source: Author's Autodesk Forma outputs, 2025.

Based on the Area metrics outputs and the concept-stage environmental analyses, all three proposals are spatially and technically feasible at this early-stage within the constraints stated in the Concept Report. The workflow therefore establishes a consistent, auditable baseline—versioned model exports and an SSOT assumption log—from which the options can proceed to financial modelling and risk/sensitivity testing in the subsequent workflow gates (Section 5.3.3). At the same time, the findings remain bounded: detailed compliance checks (e.g., code-grade daylight, structural feasibility, and final permitting documentation) are outside the scope of concept-stage modelling and must be addressed in later phases.

5.3.3 Financial Modeling, Sensitivity Analysis, and Risk Assessment

In the third empirical step, financial modeling, sensitivity analysis, and risk assessment for the three conceptual proposals were again conducted using ChatGPT 5.2 in Extended Thinking mode. Given the high complexity and data sensitivity of this step, it was structured into six sequential gates (Gate 0 to Gate 5), each corresponding to a key deliverable, ensured data transparency and auditability:

Gate workflow structure (Gate 0-Gate 5)

Before detailing each gate, an explicit overview is provided to clarify the governance logic, the frozen (SSOT-aligned) inputs, and the deliverables passed downstream. This additional layer makes the six-gate structure readable as a single auditable pipeline rather than as isolated spreadsheet tasks.

- Gate 0: Delivery directory, versioning conventions, and analysis-boundary lock-in (pre-tax, project-level KPIs).
- Gate 1: Extraction of hard constraints, program quantities, phasing, and concept deltas from the SSOT concept report.
- Gate 2: Master Inputs + Data Gaps; cost plan (WBS), schedule baseline, and cash-flow curve parameterization.
- Gate 3: Base-case financial model execution for Concepts 1-3 under a single model structure.
- Gate 4: Sensitivity (tornado) and combined market x delivery scenarios (scenario matrix).
- Gate 5: Risk register (qualitative-quantitative mapping) and Monte Carlo simulation outputs.

Each gate builds directly upon the outputs of the previous one, culminating in a consolidated deliverable package. This integrated output then supports a comprehensive comparative analysis against the earlier Concept Report, ultimately informing the real estate investment decision report.

Gate 0: Project Setup and Delivery Framework Lock-in

Method. Gate 0 establishes the delivery governance that keeps the subsequent modelling auditable: a fixed directory structure, a file naming and version-control convention, and a binding analysis boundary. In practical terms, it clarifies what is inside the feasibility test (a pre-tax, project-level appraisal) and what is outside scope (tax structuring and investment-grade financing), so that all downstream outputs remain comparable across Concepts 1–3.

Inputs. Feasibility Study – Constraints, Market & Concepts (1. Cover & Version.docx), together with the Chapter 4 governance principles applied in this empirical workflow.

Deliverables Package.

Gate0_DeliveryDirectory_Versioning_AnalysisBoundary_PreTax_SPINA3.docx (delivery directory map, file naming and version-control rules, and binding analysis-boundary policy).

Key outputs.

- Delivery directory map, including dedicated folders for Gate 0 policy decisions, evidence indexing, and a version log.
- File naming template: SPINA3_<Workstream>_<Artifact>_<Concept|All>_vX.Y_YYYYMMDD.<ext>, with major/minor increments and defined freeze points (Inputs Freeze, Model Freeze, and Final Pack).
- Binding analysis boundary: pre-tax scope and project-level KPIs (Project NPV/IRR and peak funding metrics) as the minimum comparable baseline across concepts; any non-specified conventions (e.g., time step or price basis) are marked explicitly as TBD rather than silently assumed.
- Read-only constraint checks to be enforced downstream (e.g., total GFA cap, statutory use-mix minima/maxima, public-space and infrastructure obligations, and height limits) to prevent scope drift during later model iterations.

Conclusion. By freezing governance and scope before any numerical modelling, Gate 0 reduces assumption drift and enables reproducible comparisons. The subsequent gates can therefore be read as controlled transformations from SSOT concept definitions into model-ready inputs and comparable KPIs, rather than as ad-hoc spreadsheet iteration (Gate0_DeliveryDirectory_Versioning_AnalysisBoundary_PreTax_SPINA3; Author's workflow artefact, 2025).

Gate 1: Extract developable scale and proposal parameters from the "Concept Report."

Method. Gate 1 employs ChatGPT 5.2 (Extended Thinking mode) as a controlled extractor, while Gate 1 extracts buildable capacity and scheme parameters from the Concept Report. Through the Gate 1 governance process, the Concept Report is transformed into locked, model-ready parameters: first by freezing boundary conditions and versioning rules, then by freezing constraints, program area metrics, phasing schedules, and the deltas between concept alternatives.

Inputs.

- Cover & Version.docx.
- Gate0_DeliveryDirectory_Versioning_AnalysisBoundary_PreTax_SPINA3.docx.

Deliverables Package.

- SPINA3_Inputs_HardConstraints_All_v1.0_20251224.xlsx.
- SPINA3_Inputs_ConceptProgramPhasing_C1-C3_v1.1.2_20251226.xlsx.
- SPINA3_Inputs_ConceptDeltaLog_C1-C3_v1.1_20251224.xlsx

(Author's workflow artefact, 2025).

Key outputs. A traceable set of locked constraints (Table 5.3), comparable concept program/phasing parameters (Table 5.4), and a delta log to ensure that only approved variables change across C1–C3.

Conclusion. Gate 1 transforms narrative concept material into auditable, comparable inputs for the SSOT-aligned modelling baseline established in next section.

This step corresponds to Gate 1 of the audit workflow and converts the upstream "Concept Report" into a locked, model-ready parameter set. ChatGPT 5.2 was used in extended thinking mode to extract, reconcile, and structure inputs into versioned artefacts, while preserving a strict Single Source of Truth (SSOT): the Concept Report remains the narrative baseline, and every numerical assumption used later is stored only in dedicated input spreadsheets with explicit trace references (Author's workflow artefact, 2025).

Gate 1 freezes (i) the non-negotiable planning constraints and (ii) the three concept definitions in a form that is directly usable by the financial model. The locked constraint register is treated as read-only by all downstream files, and the concept program-and-phasing register expresses each proposal as comparable quantities and delivery gates (Author's workflow artefact, 2025).

Table 5.3 Summary of the most decision-critical locked constraints (non-negotiable). Source: SPINA3_Inputs_HardConstraints_All_v1.0_20251224.xlsx, 2025.

Constraint (ID)	**Requirement / limit**	**Value**	**Unit**
Max Buildable GFA	46,006 m² gross floor area (all uses combined)	46,006	m²
Residential Use	≥40% of total GFA must be residential	≥40	% of total GFA
Tertiary/Office Use	≤40% of GFA may be tertiary (offices, "Eurotorino")	≤40	% of total GFA
Commercial/Retail	≤20% of GFA may be ASPI (retail and services)	≤20	% of total GFA
Public Service Land	≈56,649 m² (52% of site) for public spaces & facilities (parks, plazas, etc.)	≈56,649	m²
New Public Roads	Extend Via Dronero and Via Ceva through site (connect to C.so Oddone) – dedicated public right-of-way	Extend Via Dronero and Via Ceva through site; dedicated public right-of-way	text
Max building height – Via Ceva extension / internal axis	5 floors max on Via Ceva extension, internal E–W axis, and Via Savigliano; 7 floors max on other fronts (Oddone, Vigevano, etc.)	5	floors
Max building height – other fronts	5 floors max on Via Ceva extension, internal E–W axis, and Via Savigliano; 7 floors max on other fronts (Oddone, Vigevano, etc.)	7	floors
Implementation	Detailed Executive Plan or Permesso Convenzionato required (no straight permit without agreement)	Detailed Executive Plan or Permesso Convenzionato required (no straight permit without agreement)	text

These constraints are referenced explicitly in downstream quality-gate checks to ensure that no concept run exceeds the GFA cap or violates the mandated use-mix and height limits.

Table 5.4 Concept 1–3 program mix and phasing summary (as extracted from the program-and-phasing register). Source: SPINA3_Inputs_ConceptProgramPhasing_C1-C3_v1.1.2_20251226.xlsx, 2025.

Concept	Total GFA (m²)	Residential (m²)	Office / tertiary (m²)	Retail (m²)	Civic / institutional (m²)	Residential units / beds	Phasing summary
C1	46000	27600	6900	4600	6900	300 units	P1 (Years 1–2; 24 months; 150 units) → P2 (Years 3–5; 36 months; 150 units) → P3 (Year 5–6; 24 months; completion)
C2	46000	18400	18400	9200	0 (residual)	200 units	P1 (Years 1–3; 36 months; 100 units) → P2 (Years 4–6; 36 months; 100 units) → P3 (Year 6; 12 months; completion)
C3	46000	23000	13800	4600	4600	250 units; 100 beds (student housing option)	P1 (Years 1–3; 36 months; 100 units) → P2 (Years 4–5; 24 months; 150 units) → P3 (Year 5–6; 24 months; completion)

Finally, Gate 1 produces a "difference log" that records what changes across Concept 1–3 (and only those changes), so that all constant modelling rules are held fixed across runs (SPINA3_Inputs_ConceptDeltaLog_C1-C3_v1.1; Author's workflow artefact, 2025). The most decision-relevant deltas are:

- Use-mix strategy: C1 is residential-led (Residential GFA 27,600 m²), C2 approaches the statutory maximum mix (Residential and Office each 18,400 m²; Retail 9,200 m²), while C3 introduces an institutional/innovation component (Special/Institutional GFA 4,600 m²) with approx. 13,800 m² of office/tertiary space.
- Residential scale: extracted unit counts differ materially (C1 ≈300 units; C2 ≈200 units; C3 ≈250 units, with an optional student-housing component reported as 100 beds).
- Delivery rhythm: the phasing register defines concept-specific staging (Table 5.4), which later drives both peak cash need and the timing of sales/leasing receipts.
- Dependency profile: C3 explicitly depends on institutional anchoring/partnership commitments, shifting risk from pure market absorption toward governance and counterpart risk.
- Governance invariants: the same GFA cap, use-mix limits, public-space obligations and height limits apply to all concepts and are enforced through downstream constraint checks (Table 5.3).
- Traceability: every change is recorded at the level of a named input (program quantity, phasing parameter, or dependency), enabling auditable "same-model, different-inputs" comparisons in financial and risk analysis.

With these Gate 0–1 artefact locked, the workflow proceeds to establish global modelling assumptions and explicit uncertainty ranges.

Gate 2: Master Inputs and Data Gaps

Method. Gate 2 consolidates SSOT-aligned assumptions into a single Master Inputs register, while externalizes uncertainty as explicit Data Gaps. Then, defines ranges/distributions and maps costs, schedule, and receipts into time-distributed cash flows via WBS structure, baseline task timing, and curve profiles.

Inputs.

- Gate0_DeliveryDirectory_Versioning_AnalysisBoundary_PreTax_SPINA3.docx.
- SPINA3_Inputs_HardConstraints_All_v1.0_20251224.xlsx.
- SPINA3_Inputs_ConceptProgramPhasing_C1-C3_v1.1.2_20251226.xlsx.
- SPINA3_Inputs_ConceptDeltaLog_C1-C3_v1.1_20251224.xlsx

Deliverables Package.

- SPINA3_Inputs_MasterInputs_All_v1.4.2.1_C2GFAFix_TraceHotfix_20251227.xlsx.
- SPINA3_Inputs_DataGapsRanges_All_v1.3_20251227.xlsx.
- SPINA3_Hold_Operations_AssumptionsPack_ScenarioB_v1.0.1.1_20251227_VersionLogAdd.xlsx.
- SPINA3_CostPlan_WBS_C1-C3_v1.2.1_20251226.xlsx.
- SPINA3_ScheduleBaseline_Gantt_C1-C3_v1.0.4.5_20251227_UltraTraceRefPatch.xlsx.
- SPINA3_CashFlowCurves_Params_C1-C3_v1.0.1.4_20251227_UltraTraceRefPatch.xlsx

(Author's workflow artefact, 2025).

Key outputs. A model-ready baseline of deterministic inputs (Table 5.5), quantified uncertainties for sensitivity testing (Table 5.6), and a reconciled cost-and-schedule baseline suitable for comparability across concepts (Table 5.7).

Conclusion. Together, Gates 2 define the controlled 'input envelope' that allows financial model to compute finance KPIs consistently and sensitivity testing to test downside exposure.

Gate 2 translate the locked concept quantities into a comparable modelling baseline by (i) consolidating all shared assumptions into a single global table ("Master Inputs") and (ii) making all remaining uncertainties explicit through a formal Data Gaps register with ranges/distributions. This two-layer structure is a governance mechanism: Master Inputs prevent silent assumption drift across concept runs, while Data Gaps ensure that uncertainty is not "hidden" inside the model but is represented as declared variables that can be stress-tested later (Author's workflow artefact, 2025).

Within this setup, four artefact groups define the modelling boundary and the translation from static rates to time-distributed cash flows: (1) the Hold & Operations pack defines leasing/occupancy/free-rent and exit assumptions for any held components; (2) the Cost Plan WBS establishes the cost structure and totals by category; (3) the Schedule baseline defines task timing and phase overlap; and (4) the Cash-flow curves parameterize how each cost and receipt stream is distributed over time. All four are versioned and trace-linked so that each Gate 4 model run can be reproduced from a frozen input set (Author's workflow artefact, 2025).

Table 5.5 Master Inputs snapshot (selected key assumptions used downstream). Source: SPINA3_Inputs_MasterInputs_All_v1.4.2.1_C2GFAFix_TraceHotfix_20251227.xlsx, 2025.

Assumption (ID)	**Base value**	**Unit**	**Source location**
MI-0129 — Residential sales price & rent assumptions (new apartments)	2050	€/m²	MasterInputs, Master_Inputs!A130:Q130
MI-HOLD-0101 — Res_Rent_EUR_m2_month	7.5	€/m²/month	MasterInputs, Master_Inputs!A193:Q193
MI-HOLD-0201 — Office_Rent_EUR_m2_year	100	€/m²/year	MasterInputs, Master_Inputs!A203:Q203
MI-HOLD-0301 — Retail_Rent_EUR_m2_year	120	€/m²/year	MasterInputs, Master_Inputs!A213:Q213
MI-HOLD-0102 — Res_LeaseUp	8	units/month	MasterInputs, Master_Inputs!A194:Q194
MI-HOLD-0202 — Office_LeaseUp	300	m²/month	MasterInputs, Master_Inputs!A204:Q204
MI-HOLD-0302 — Retail_LeaseUp	150	m²/month	MasterInputs, Master_Inputs!A214:Q214
MI-HOLD-0104 — Res_Vacancy	5	%	MasterInputs, Master_Inputs!A196:Q196
MI-HOLD-0204 — Office_Vacancy	8	%	MasterInputs, Master_Inputs!A206:Q206
MI-HOLD-0304 — Retail_Vacancy	10	%	MasterInputs, Master_Inputs!A216:Q216
MI-HOLD-0205 — Office_FreeRent_Months	6	months	MasterInputs, Master_Inputs!A207:Q207
MI-HOLD-0305 — Retail_FreeRent_Months	6	months	MasterInputs, Master_Inputs!A217:Q217
MI-HOLD-0002 — Hold_Share_Residential	20	%	MasterInputs, Master_Inputs!A190:Q190
MI-HOLD-0003 — Hold_Share_Office	100	%	MasterInputs, Master_Inputs!A191:Q191
MI-HOLD-0110 — Res_Exit_CapRate_or_Yield	5.1	%	MasterInputs, Master_Inputs!A202:Q202
MI-HOLD-0210 — Office_Exit_CapRate_or_Yield	7.6	%	MasterInputs, Master_Inputs!A212:Q212
MI-HOLD-0310 — Retail_Exit_CapRate_or_Yield	8.38	%	MasterInputs, Master_Inputs!A222:Q222
MI-HOLD-0109 — Res_HoldPeriod_Years	10	years	MasterInputs, Master_Inputs!A201:Q201
DG-0052 — NPV discount rate (pre-tax, nominal)	11.0	%/year	DataGapsRanges, Ranges_ModelReady!A2:O2
DG-0002 — Inflation / general escalation (nominal p.a.)	2.0	%/year	DataGapsRanges, Ranges_ModelReady!A3:O3

Where an assumption is classified as a Data Gap, its base value is still referenced by the deterministic base-case run, but its range/distribution is preserved for scenario and sensitivity testing (Author's workflow artefact, 2025).

Table 5.6 Data gaps and adopted ranges/distributions (model-ready selection). Source: SPINA3_Inputs_DataGapsRanges_All_v1.3_20251227.xlsx, 2025.

Gap_ID	Uncertainty / parameter	Base	Low	High	Distribution	Affects module
DG-0052	NPV discount rate (pre-tax, nominal)	11.0	9.0	14.0	Triangular	MasterInputs; FinanceModel
DG-0002	Inflation / general escalation (nominal p.a.)	2.0	1.0	3.0	Triangular	MasterInputs; FinanceModel; HoldPack (proxy)
DG-0059	Cost escalation (nominal p.a., applied to Hard+Soft costs)	2.0	1.0	3.0	Triangular	CostPlan; FinanceModel
DG-0060	Rent escalation / indexation (nominal p.a.)	2.0	1.0	3.0	Triangular	HoldPack; FinanceModel
DG-0061	Approvals duration (months)	6.0	4.0	12.0	Triangular	ScheduleBaseline; CashFlowCurves (timing)
DG-0062	Presales / absorption duration (months)	12.0	9.0	18.0	Triangular	ScheduleBaseline; CashFlowCurves (receipts)
DG-0041	Construction contingency (% of HardCost)	0.1	0.07	0.15	Triangular	CostPlan; FinanceModel
DG-0042	Schedule contingency / delay buffer (months)	3.0	0.0	12.0	Triangular	ScheduleBaseline; FinanceModel
DG-0010	Schedule duration multiplier (applies to Assumption_Inputs for C1)	1.0	0.9	1.2	Triangular	ScheduleBaseline; CashFlowCurves
DG-0014	Schedule duration multiplier (applies to Assumption_Inputs for C2)	1.0	0.9	1.2	Triangular	ScheduleBaseline; CashFlowCurves
DG-0019	Schedule duration multiplier (applies to Assumption_Inputs for C3)	1.0	0.9	1.2	Triangular	ScheduleBaseline; CashFlowCurves
DG-0048	Residential hold share (%)	20.0	10.0	30.0	Triangular	MasterInputs; HoldPack; FinanceModel
DG-0054	Residential market rent (€/m²/month) for hold underwriting	7.5	6.0	9.0	Triangular	HoldPack; FinanceModel; CashFlowCurves
DG-0022	Total cost uncertainty multiplier (derived from CostPlan Low/High totals)	1.0	0.811	1.189	Triangular	CostPlan; CashFlowCurves; FinanceModel

Table 5.7 Baseline cost and schedule summary (WBS totals, timing, and applied cash-flow curve profiles). Source: SPINA3_CostPlan_WBS_C1-C3_v1.2.1_20251226.xlsx, 2025; SPINA3_ScheduleBaseline_Gantt_C1-C3_v1.0.4.5_20251227_UltraTraceRefPatch.xlsx, 2025.

Concept	Land / site (€ m)	Hard costs (€ m)	Public works (€ m)	Soft costs (€ m)	Total base cost (€ m)	Baseline duration (months)	Phase windows (months from start)	Cash-flow curve profiles
C1	5.9	72.5	4.0	11.5	93.9	87	P1:0–54; P2:9–75; P3:33–87	Spend: SPEND_12_C1 (12m); Receipt: RECEIPT_12_C1 (12m)
C2	5.9	72.6	4.0	11.6	94.1	87	P1:0–66; P2:21–87; P3:45–87	Spend: SPEND_12_C2 (12m); Receipt: RECEIPT_12_C2 (12m)
C3	6.1	74.0	4.0	11.8	95.9	87	P1:0–66; P2:21–75; P3:33–87	Spend: SPEND_12_C3 (12m); Receipt: RECEIPT_12_C3 (12m)

In Table 5.7, cost totals are aggregated from the detailed WBS line items (CostPlan_WBS; Author's workflow artefact, 2025), while timing is taken from the baseline task register (ScheduleBaseline_Gantt; Author's workflow artefact, 2025). The cash-flow curve profiles used to distribute each stream are referenced by ID in the curve library (CashFlowCurves_Params; Author's workflow artefact, 2025). Together, these Gate 2 artefacts create a consistent input envelope for the financial model.

Gate 3: Financial Model for Conceptual Proposals

Method. Gate 3, a single pre-tax, project-level model structure is executed three times (C1–C3) under identical SSOT-aligned assumptions (discounting, escalation, cost rules, and operational boundary), allowing only the concept delta variables and phasing inputs to differ.

Inputs.

- SPINA3_Inputs_MasterInputs_All_v1.4.2.1_C2GFAFix_TraceHotfix_20251227.xlsx.
- SPINA3_Inputs_DataGapsRanges_All_v1.3_20251227.xlsx.
- SPINA3_Hold_Operations_AssumptionsPack_ScenarioB_v1.0.1.1_20251227_VersionLogAdd.xlsx.
- SPINA3_CostPlan_WBS_C1-C3_v1.2.1_20251226.xlsx.
- SPINA3_ScheduleBaseline_Gantt_C1-C3_v1.0.4.5_20251227_UltraTraceRefPatch.xlsx.
- SPINA3_CashFlowCurves_Params_C1-C3_v1.0.1.4_20251227_UltraTraceRefPatch.xlsx

Deliverables Package.

SPINA3_FinancialModel_C1-C3_v1.0_20251227.xlsx (Author's workflow artefact, 2025).

Key outputs. Headline base-case KPIs and built-in constraint/sanity checks that validate input consistency (Table 5.8; Constraints Checks).

Conclusion. Gate 3 produces directly comparable financial outputs that can be interpreted through sensitivity and scenario testing, rather than treated as standalone point estimates.

Gate 3 implements the financial model as a "same model, three runs" structure: a single spreadsheet model architecture is executed for Concept 1–3 using the locked constraints, shared Master Inputs, declared Data Gaps (base values), and the concept-specific program/phasing deltas. This design is intentional: it reduces model-structure bias and ensures that any difference in outputs is attributable to the concept inputs rather than to inconsistent modelling logic (Author's workflow artefact, 2025).

Comparability controls are enforced by freezing: (i) the analysis boundary (pre-tax, project-level KPIs), (ii) the market base assumptions and hold/operations rules, (iii) the cost structure and unit-rate logic, and (iv) the schedule baseline and cash-flow curve mapping. Only the concept deltas (use mix, quantities, and phasing) are allowed to vary across C1–C3 (Author's workflow artefact, 2025).

Table 5.8 Base-case KPI outputs for Concept 1–3 (Gate 3 financial model; pre-tax, project-level). Source: SPINA3_FinancialModel_C1-C3_v1.0_20251227.xlsx, 2025.

KPI	Unit	Concept 1	Concept 2	Concept 3	Source (sheet/cells)
Project NPV (€)	€	-34 449 181	-37 033 242	-37 100 988	Outputs_KPIs!A4:D4
Project IRR (annual)	%	-17.3%	-19.8%	-22.4%	Outputs_KPIs!A5:D5
Peak Cash Need (€)	€	71 251 128	75 499 897	90 235 961	Outputs_KPIs!A8:D8
Peak Equity Need (€)	€	46 549 385	54 505 546	51 429 092	Outputs_KPIs!A11:D11
Equity Multiple	x	0	0	0	Outputs_KPIs!A10:D10

The model embeds dedicated constraint-check statements to confirm that each run respects the locked boundary conditions: total GFA remains within the HC-003 cap, use-mix share checks (HC-005 to HC-007) pass, and phase allocations reconcile to the concept total (Constraints_Checks; Author's workflow artefact, 2025). As a cost-structure indicator, public works are reported as approximately 4.2% of total base cost for all three concepts (Constraints_Checks; Author's workflow artefact, 2025).

Under the base-case assumptions captured in Gate 1, all three concepts produce negative project-level NPVs and negative IRRs (Table 5.8), indicating that the current concept package does not yet clear the required feasibility thresholds under a pre-tax project return test. Among the three, Concept 1 is the least unfavorable on Project NPV (−34.4 M€) and has the lowest peak cash need (71.3 M€), consistent with its residential-led program and earlier sell-through logic. Concept 2 and Concept 3 show more negative NPVs (−37.0 M€ and −37.1 M€ respectively) and higher peak cash requirements, reflecting larger non-residential components and longer delivery and stabilisation dynamics implied by their program mixes and phasing profiles (Tables 5.4 and 5.7; Author's workflow artefact, 2025).

These outputs remain concept-gate and assumption-driven. They are sensitive to the declared uncertainty ranges (Table 5.6), and they do not substitute for a detailed design-development cost estimate, confirmed financing terms, or binding market commitments. Accordingly, Gate 4 is used to identify which assumptions dominate value creation and downside exposure before any recommendation is made.

Gate 4: Sensitivity and Scenario Testing

Method. Gate 4. Identical variable shocks are applied across concepts to (i) rank the most influential drivers of Project NPV/IRR (tornado) and (ii) construct a 3×3 market × delivery scenario matrix to test downside resilience and upside potential.

Inputs.

- SPINA3_FinancialModel_C1-C3_v1.0_20251227.xlsx
- SPINA3_Inputs_MasterInputs_All_v1.4.2.1_C2GFAFix_TraceHotfix_20251227.xlsx.
- SPINA3_Inputs_DataGapsRanges_All_v1.3_20251227.xlsx.
- SPINA3_Hold_Operations_AssumptionsPack_ScenarioB_v1.0.1.1_20251227_VersionLogAdd.xlsx.

Deliverables Package.

- SPINA3_Sensitivity_Tornado_v1.2.2_20251227_2pagesPerConcept_CleanFilled.pdf
- SPINA3_Scenario_Matrix_v1.0.1_20251227_Viz_v1.xlsx

(Author's workflow artefact, 2025).

Key outputs. Embedded tornado charts (Figures 5.11–5.13), a scenario matrix visual (Figure 5.13), and comparable summary tables for decision drivers and scenario outcomes (Tables 5.9–5.10).

Conclusion. Gate 4 identifies which assumptions dominate value creation and which downside conditions would trigger a re-check of market inputs, cost/schedule buffers, or concept phasing.

Gate 4 stress-tests the Gate 3 base case using identical variables across concepts to identify decision drivers and to characterize downside exposure. Two complementary instruments are used: (i) tornado sensitivity charts, which isolate the marginal impact of key variables on project KPIs, and (ii) a structured 3×3 scenario matrix, which combines market and delivery/capital conditions into comparable downside/base/upside cases (Author's workflow artefact, 2025).

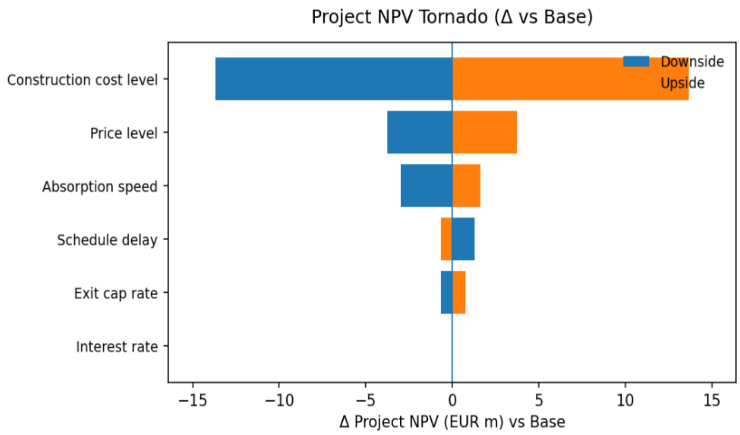

Figure 5.13 Tornado sensitivity chart – Concept 1 (Gate 4 output). Source: SPINA3_Sensitivity_Tornado_v1.2.2_20251227_2pagesPerConcept_CleanFilled.pdf, 2025 (pp. 1–6).

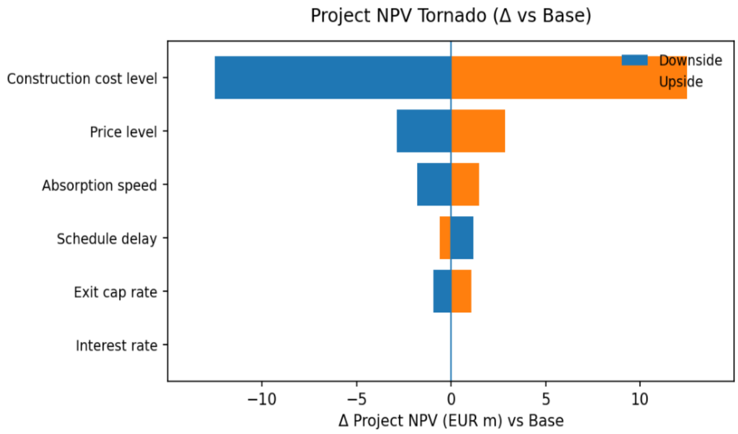

Figure 5.14 Tornado sensitivity chart – Concept 2 (Gate 4 output). Source: SPINA3_Sensitivity_Tornado_v1.2.2_20251227_2pagesPerConcept_CleanFilled.pdf, 2025 (pp. 1–6).

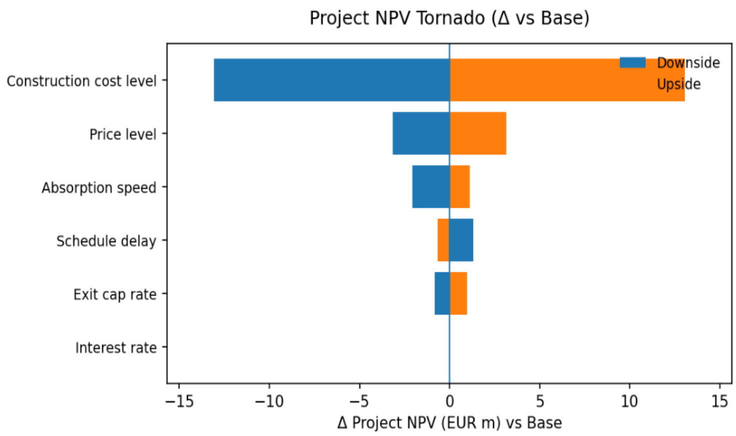

Figure 5.15 Tornado sensitivity chart – Concept 3 (Gate 4 output). Source: SPINA3_Sensitivity_Tornado_v1.2.2_20251227_2pagesPerConcept_CleanFilled.pdf, 2025 (pp. 1–6).

Table 5.9 Top sensitivity drivers (Tornado; ranked by visual ordering of the Gate 4 output). Source: SPINA3_Sensitivity_Tornado_v1.2.2_20251227_2pagesPerConcept_CleanFilled.pdf, 2025.

Concept	Rank	Driver (Tornado)
C1	1	Price level
C1	2	Construction cost level
C1	3	Absorption speed
C1	4	Schedule delay
C1	5	Exit cap rate
C2	1	Price level
C2	2	Construction cost level
C2	3	Absorption speed
C2	4	Schedule delay
C2	5	Exit cap rate
C3	1	Price level
C3	2	Construction cost level
C3	3	Absorption speed
C3	4	Schedule delay
C3	5	Exit cap rate

Across all three concepts, the tornado charts indicate a consistent ordering of dominant drivers: price level and construction cost level are the most influential, followed by absorption speed, schedule delay, and the exit cap rate (Table 5.9). The directionality is as expected for concept-gate feasibility: higher achieved prices/rents improve NPV and IRR; higher construction costs reduce NPV and worsen IRR; faster absorption improves returns by accelerating receipts and reducing carry; schedule delays worsen returns by pushing costs earlier and receipts later; and higher exit cap rates reduce terminal value (Author's workflow artefact, 2025).

Table 5.10 Scenario matrix results (diagonal cases: Upside_M/Upside_D, Base_M/Base_D, Downside_M/Downside_D). Source: SPINA3_Scenario_Matrix_v1.0.1_20251227_Viz_v1.xlsx, 2025.

Concept	Scenario	Project NPV (EUR m)	Project IRR (annual %)	Peak Cash Need (EUR m)
C1	Upside	-14.3	-0.8%	62.0
C1	Base	-25.2	-7.6%	71.3
C1	Downside	-34.2	-10.9%	86.0
C2	Upside	-13.5	0.8%	66.1
C2	Base	-22.6	-4.8%	75.5
C2	Downside	-30.5	-7.9%	88.9
C3	Upside	-15.6	-2.3%	77.8
C3	Base	-25.2	-9.1%	90.2
C3	Downside	-32.6	-13.0%	102.6

3×3 Scenario Matrix (Market × Delivery): Project NPV

C1 – Project NPV (EUR m)

	Upside_D	Base_D	Downside_D
Upside_M	-14.3	-21.9	-26.7
Base_M	-17.8	-25.2	-29.7
Downside_M	-23.0	-30.2	-34.2

C2 – Project NPV (EUR m)

	Upside_D	Base_D	Downside_D
Upside_M	-13.5	-20.7	-25.4
Base_M	-15.4	-22.6	-27.1
Downside_M	-19.6	-26.6	-30.5

C3 – Project NPV (EUR m)

	Upside_D	Base_D	Downside_D
Upside_M	-15.6	-22.9	-27.5
Base_M	-17.9	-25.2	-29.6
Downside_M	-21.5	-28.6	-32.6

Figure 5.16 3×3 scenario matrix visualization (Market × Delivery shocks): Project NPV (EUR m) for Concepts 1–3 (Gate 4 output; derived from SPINA3_Scenario_Matrix_v1.0.1_20251227_Viz_v1.xlsx, sheets Matrix_C1–Matrix_C3). Source: SPINA3_Scenario_Matrix_v1.0.1_20251227_Viz_v1.xlsx, 2025 (sheets Matrix_C1–Matrix_C3).

For decision-making, Gate 4 clarifies that (i) value creation is primarily governed by achievable pricing/rents and construction cost control, and (ii) delivery risk (schedule slippage and absorption slowdown) materially increases peak cash need and deepens negative NPV in downside cases (Table 5.10). In this baseline, even the upside diagonal cases remain negative on project NPV across concepts, suggesting that feasibility would require either a different program/tenure structure, material cost reductions, improved market pricing, or policy/partner interventions to shift the value envelope. These findings provide the basis for the subsequent discussion on concept selection and on which assumptions warrant revisiting before advancing to a more detailed design and financing phase.

Gate 5: Risk Assessment

Method. In Gate 5, risks were first structured in a register format using consistent categories and a 1–5 probability score and a 1–5 impact score, with severity defined as the product of the two scores (SPINA3_RiskRegister_v1.0.4.xlsx, README sheet). Each risk row records its trace reference to the SSOT artefacts and, where applicable, indicates whether it can be represented as a model variable for uncertainty testing (SPINA3_RiskRegister_v1.0.4.xlsx, RiskRegister sheet). The same uncertainty set is then tested probabilistically using Monte Carlo simulation for Concepts 1–3, reporting distribution statistics and threshold probabilities (MonteCarlo_Results_C1-C3_v1.0.2.xlsx, Summary sheet).

Inputs.

- SPINA3_FinancialModel_C1-C3_v1.0_20251227.xlsx.
- SPINA3_Inputs_MasterInputs_All_v1.4.2.1_C2GFAFix_TraceHotfix_20251227.xlsx.
- SPINA3_Inputs_DataGapsRanges_All_v1.3_20251227.xlsx.
- SPINA3_Hold_Operations_AssumptionsPack_ScenarioB_v1.0.1.1_20251227_VersionLogAdd.xlsx.
- SPINA3_CostPlan_WBS_C1-C3_v1.2.1_20251226.xlsx.
- SPINA3_ScheduleBaseline_Gantt_C1-C3_v1.0.4.5_20251227_UltraTraceRefPatch.xlsx.
- SPINA3_CashFlowCurves_Params_C1-C3_v1.0.1.4_20251227_UltraTraceRefPatch.xlsx.
- SPINA3_Sensitivity_Tornado_v1.2.2_20251227_2pagesPerConcept_CleanFilled.pdf.
- SPINA3_Scenario_Matrix_v1.0.1_20251227_Viz_v1.xlsx.

Deliverables Package.

- SPINA3_RiskRegister_v1.0.4_20251227.xlsx.
- SPINA3_MonteCarlo_Results_C1-C3_v1.0.2_20251227.xlsx.

(Author's workflow artefacts, 2025).

Key outputs. The risk register contains 76 risks (RR-001 to RR-076). The most populated categories are Schedule (17), Planning / Policy (16), Operations & Exit (14), Financing (12), Cost (9), Market (6), and Environment & Geotechnical (2). Table 5.11 lists the top-ranked risks by severity score as defined in the workbook.

Conclusion. Gate 5 consolidates risk reasoning into a single, traceable structure and then tests the same uncertainty set probabilistically across Concepts 1–3. In built-environment feasibility terms, it connects qualitative project risks (planning, market, cost, schedule, operations) to quantifiable variables and thereby evaluates robustness, not only point-estimate performance.

Table 5.11 Top 10 risks by severity score (ID, category, probability/impact, mitigation/response). Source: SPINA3_RiskRegister_v1.0.4_20251227.xlsx, RiskRegister sheet (rows 2–77; columns A–N; scoring fields in columns I, K, L).

Risk_ID	Risk_Title	Category	Prob (1–5)	Impact (1–5)	Severity (P×I)	Mitigation	Response
RR-002	Nominal vs real terms; inflation/escalation policy	Financing	4	5	20	Decide nominal vs real convention and escalation policy; document in Master Inputs as app…	Allow contingency / scenario testing; refine inputs when evidence becomes available.
RR-047	Residential tenure split (for-sale vs hold-to-rent)	Market	4	5	20	Decide tenure split policy (for-sale vs hold-to-rent) for base case and sensitivities; re…	Allow contingency / scenario testing; refine inputs when evidence becomes available.
RR-048	Hold period (years) for rental component	Operations & Exit	4	5	20	Assume hold period and exit method for rental component as scenario; validate later.	Allow contingency / scenario testing; refine inputs when evidence becomes available.
RR-051	NPV discount rate basis (pre-tax)	Financing	4	5	20	Set pre-tax discount rate basis as policy; include in Outputs as assumption overlay with …	Allow contingency / scenario testing; refine inputs when evidence becomes available.
RR-059	Rent escalation / indexation (nominal p.a.)	Market	4	5	20	Use HoldPack rent escalation proxy (ECB inflation aim) until lease indexation known.	Allow contingency / scenario testing; refine inputs when evidence becomes available.
RR-072	Vacancy / credit loss risk (held components)	Operations & Exit	4	5	20	Confirm and update assumptions with primary market / technical evidence; adjust phasing a…	Maintain flexibility (phasing, optionality); include contingency and downside scenario.
RR-073	Operating cost risk (non-recoverable	Operations & Exit	4	5	20	Confirm and update assumptions with primary market /	Maintain flexibility (phasing, optionality); include

Risk_ID	Risk_Title	Category	Prob (1–5)	Impact (1–5)	Severity (P×I)	Mitigation	Response
	opex % of GRI)					technical evidence; adjust phasing a…	contingency and downside scenario.
RR-075	Hold-share / exit strategy policy risk	Operations & Exit	4	5	20	Confirm and update assumptions with primary market / technical evidence; adjust phasing a…	Maintain flexibility (phasing, optionality); include contingency and downside scenario.
RR-030	Design & engineering fees basis	Cost	4	4	16	Provide cost basis and breakdown; until available, assume with scenarios and record as 'A…	Allow contingency / scenario testing; refine inputs when evidence becomes available.
RR-031	Permitting/ authority fees basis	Cost	4	4	16	Provide cost basis and breakdown; until available, assume with scenarios and record as 'A…	Allow contingency / scenario testing; refine inputs when evidence becomes available.

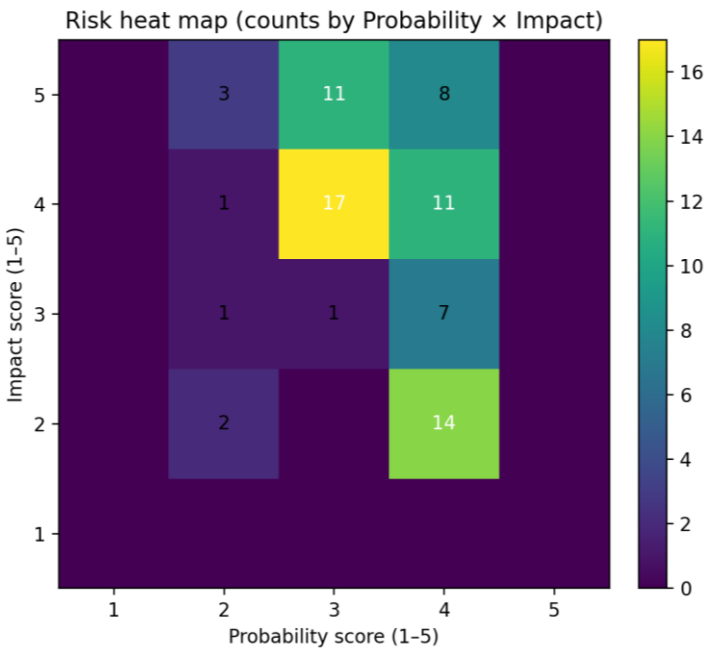

Figure 5.17 Risk heat map (Probability × Impact; counts by score pair). Source: SPINA3_RiskRegister_v1.0.4_20251227.xlsx, RiskRegister sheet (Probability_Score_1to5, Impact_Score_1to5; derived count matrix).

For probabilistic feasibility, the Monte Carlo summaries indicate that the simulated distributions for all three concepts remain on the negative side of the NPV threshold under the current uncertainty set, while peak funding gaps remain substantial (MonteCarlo_Results_C1-C3_v1.0.2.xlsx, Summary sheet). Table 5.12 reports the available percentile statistics and threshold probabilities.

Table 5.12 Monte Carlo summary statistics by concept (P10/P50/P90 for NPV and IRR; threshold probabilities; peak funding gap). Source: SPINA3_MonteCarlo_Results_C1-C3_v1.0.2_20251227.xlsx, Summary sheet (rows 2–4; columns A–AA).

Concept	NPV P10	NPV P50	NPV P90	IRR P10	IRR P50	IRR P90	Prob (NPV >0)	Prob (IRR>11%)	Peak funding gap P50
C1	€-42,589,556	€-34,878,376	€-27,184,898	-21.8%	-17.5%	-13.0%	0.0%	0.0%	€72,305,222
C2	€-44,224,428	€-37,128,217	€-30,117,235	-24.1%	-20.0%	-15.9%	0.0%	0.0%	€76,901,525
C3	€-44,164,184	€-36,809,034	€-29,760,224	-27.1%	-22.4%	-17.8%	0.0%	0.0%	€90,674,746

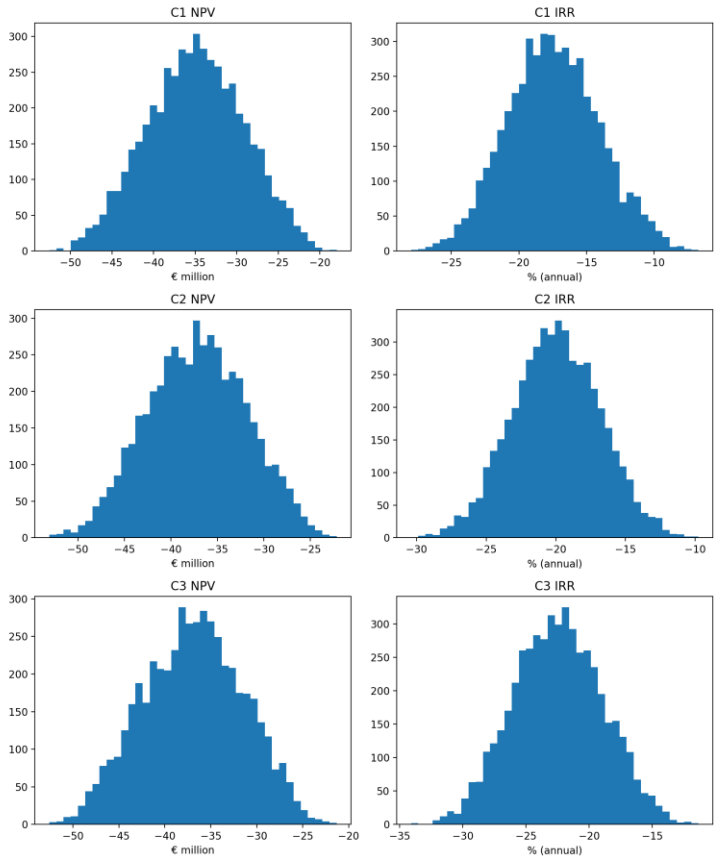

Figure 5.18 Monte Carlo distributions (NPV and IRR histograms) for C1–C3. Source: SPINA3_MonteCarlo_Results_C1-C3_v1.0.2_20251227.xlsx, Sim_C1/Sim_C2/Sim_C3 (NPV, IRR columns).

The integrated risk analysis demonstrates that, within the SSOT assumptions and the uncertainty ranges tested, concept choice must be justified not only by deterministic KPIs but by robustness under downside variation. The outputs also clarify which assumption families (market absorption, exit/hold strategy, discounting conventions, and operating performance) dominate feasibility, thereby setting explicit conditions to be stated in the consolidated decision report.

5.3.4 Integrated Comparative Decision-Making

Method. The third empirical step aimed at synthesizing the findings from prior gates and integrating them into a real estate investment decision report using ChatGPT-5.2's Deep Research mode. The report is generated by summarizing (i) non-negotiable constraints and program definitions (SSOT), (ii) deterministic feasibility outputs and their main drivers, and (iii) downside exposure identified through sensitivity, scenario testing, and Monte Carlo distributions. The decision stance is written as conditional, meaning that feasibility is framed through explicit enabling conditions (e.g., delivery sequencing and public/private risk-sharing), rather than presented as an unconditional 'best' concept.

Inputs.

- SPINA3_Inputs_HardConstraints_All_v1.0_20251224.xlsx.
- SPINA3_Inputs_ConceptProgramPhasing_C1-C3_v1.1.2_20251226.xlsx.
- SPINA3_FinancialModel_C1-C3_v1.0_20251227.xlsx.
- SPINA3_Inputs_MasterInputs_All_v1.4.2.1_C2GFAFix_TraceHotfix_20251227.xlsx.
- SPINA3_Inputs_DataGapsRanges_All_v1.3_20251227.xlsx.
- SPINA3_Hold_Operations_AssumptionsPack_ScenarioB_v1.0.1.1_20251227_VersionLogAdd.xlsx.
- SPINA3_CostPlan_WBS_C1-C3_v1.2.1_20251226.xlsx.
- SPINA3_ScheduleBaseline_Gantt_C1-C3_v1.0.4.5_20251227_UltraTraceRefPatch.xlsx.
- SPINA3_CashFlowCurves_Params_C1-C3_v1.0.1.4_20251227_UltraTraceRefPatch.xlsx.
- SPINA3_Sensitivity_Tornado_v1.2.2_20251227_2pagesPerConcept_CleanFilled.pdf.
- SPINA3_Scenario_Matrix_v1.0.1_20251227_Viz_v1.xlsx.
- SPINA3_RiskRegister_v1.0.4_20251227.xlsx.
- SPINA3_MonteCarlo_Results_C1-C3_v1.0.2_20251227.xlsx.

Deliverables Package.

Integrated Investment Decision Report – SPINA 3 Redevelopment (Turin, Italy) – Revised_v1.1.docx (Author's workflow artefacts, 2025).

Key outputs. The decision report adopts a conditional recommendation: it prioritizes a housing-led approach akin to Concept 1, while stating that Concepts 2 and 3 are not investable under current assumptions without additional enabling measures or partner commitments (Integrated Investment Decision Report – Revised_v1.1.docx, Section 1 and Section 7). Table 5.13 summarizes the report's comparative stance and the conditions/mitigations explicitly attached to each concept option.

Conclusion. This report does not introduce a new model layer; instead, it synthesizes the previously established constraints, base-case feasibility outputs, sensitivity/scenario findings, and Monte Carlo robustness results in an explicit decision stance suitable for design development and stakeholder discussion.

Table 5.13 Decision summary matrix (Concept 1–3): feasibility stance, strengths, binding risks, conditions/mitigations, and decision implication. Source: Integrated Investment Decision Report – SPINA 3 Redevelopment (Turin, Italy) – Revised_v1.1.docx, Sections 1, 7–8 (synthesized).

Concept	Feasibility stance (report)	Key strengths (qualitative)	Binding risks	Required conditions / mitigations	Decision implication
C1	Conditionally recommended	Aligns with housing demand; simpler program; resilience to non-leasing components; stronger public acceptance (park + housing).	Project-level feasibility remains negative under base assumptions; dependency on public-sector support and delivery phasing discipline.	Secure public cost-sharing/incentives and an enabling land arrangement; phase optional non-residential components after housing delivery; adopt go/no-go milestones.	Proceed only under defined public–private risk-sharing conditions; otherwise defer.
C2	Not recommended under current assumptions	Potential to catalyze a mixed-use hub if office/retail perform strongly.	High exposure to tertiary/retail demand and absorption; downside sensitivity makes feasibility unlikely without substantial external support.	Would require exceptional market traction and significant subsidies; not prioritized for advancement.	Defer / exclude from the preferred concept set.
C3	Conditionally possible only with partner-funded anchors	Strategic alignment with innovation ecosystem; potential long-term place-making value.	Large financial gap under base assumptions; relies on pre-commitment from institutional partners and public enabling actions.	Secure institutional anchor tenants / partner financing and link to public programs; otherwise treat innovation components as optional add-ons.	Advance only if partner commitments are obtained; otherwise do not pursue as standalone.

This integrated comparative gate translates quantitative outputs into an auditable decision narrative: it clarifies what must be true (conditions and mitigations) for the preferred concept to be advanced and documents why other concepts are deprioritized under the same evidence base.

5.3.5 Supplement with Visual Charts

Method. In the fifth step, Nano Banana Pro was used to standardize chart layout and typography, consolidate multi-concept outputs into single-page comparative figures, and correct common presentation issues (misaligned axes, inconsistent labels, and mixed units). The refinements focus on: (i) keeping identical scales across concepts; (ii) preserving the original ordering/definitions of variables; and (iii) maintaining explicit source trace references in figure captions.

Key outputs. The refined figure set includes: (1) vertically stacked multi-concept schedules; (2) aligned cash-flow curves; (3) publication-ready tornado; (4) scenario visuals; and (5) consistently formatted

Monte Carlo plots with percentile markers. Table 5.14 summarizes the refined figure types and their traceable data sources.

Conclusion. This step improves the communicative quality of the analytical outputs for thesis readability and stakeholder-style comparison. It is treated as a post-processing step: visual clarity is enhanced without changing any underlying numeric values, which remain traceable to the SSOT artefacts.

Table 5.14 Visual refinement log (chart type, refinement actions, and traceable data source). Source: Author's post-processing gate (2025); underlying data sources listed per row.

Figure / chart type	Refinement actions (Nano Banana Pro)	Traceable data source
Multi-concept Gantt comparison	Standardize time axis and typography; stack concepts vertically for direct reading; keep phase labels consistent.	SPINA3_ScheduleBaseline_Gantt_C1-C3_v1.0.4.5_20251227_UltraTraceRefPatch.xlsx (baseline schedule).
Cash-flow curves (C1–C3)	Plot all concepts on common scale; align legends and units; remove redundant gridlines; ensure identical start date.	SPINA3_CashFlowCurves_Params_C1-C3_v1.0.1.4_20251227_UltraTraceRefPatch.xlsx (curve parameters).
Tornado sensitivity charts	Harmonize variable naming; align bar order and axis limits across concepts; export publication-ready images.	SPINA3_Sensitivity_Tornado_v1.2.2_20251227_2pagesPerConcept_CleanFilled.pdf (concept pages).
Scenario matrix visualization (3×3)	Correct label placement; ensure consistent color meaning; consolidate C1–C3 into a single page set.	SPINA3_Scenario_Matrix_v1.0.1_20251227_Viz_v1.xlsx (Matrix_C1–Matrix_C3).
Monte Carlo distribution and CDF plots	Convert model outputs to consistent €/% units; unify binning and axes; add percentile markers (P10/P50/P90) without altering values.	SPINA3_MonteCarlo_Results_C1-C3_v1.0.2_20251227.xlsx (Summary).

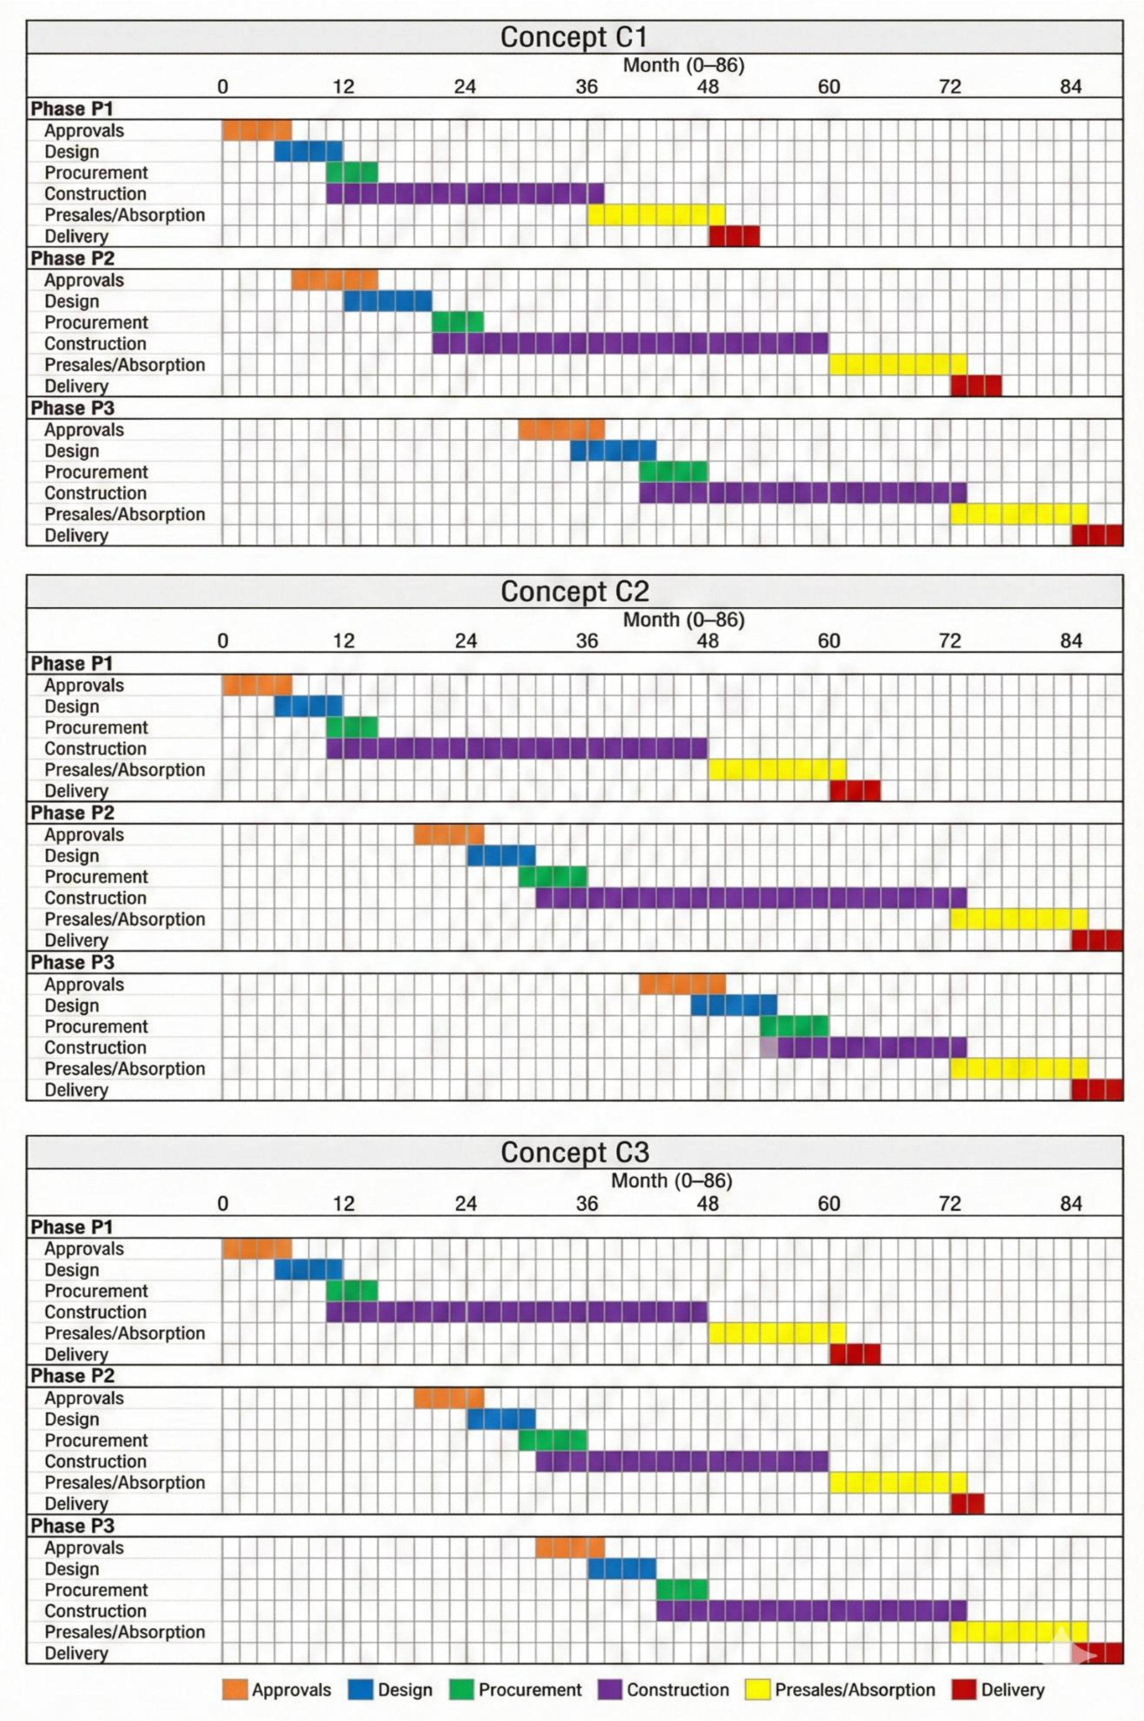

Figure 5.19 Multi-concept Gantt comparison. Source: Author's Nano Banana Pro outputs, 2025.

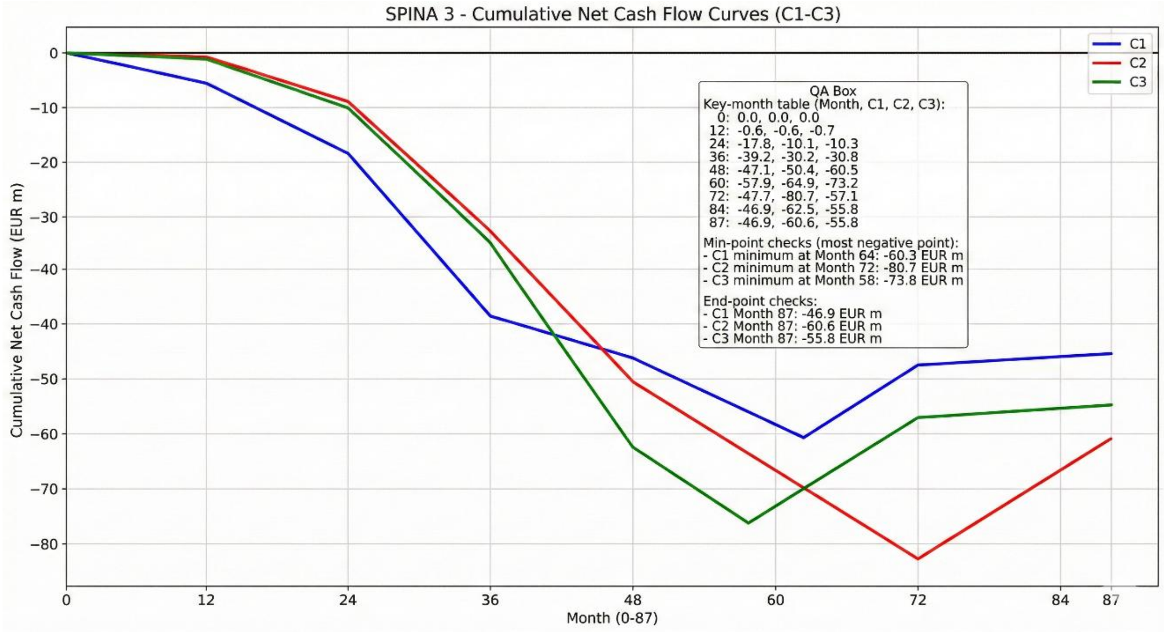

Figure 5.20 Cash-flow curves (C1–C3). Source: Author's Nano Banana Pro outputs, 2025.

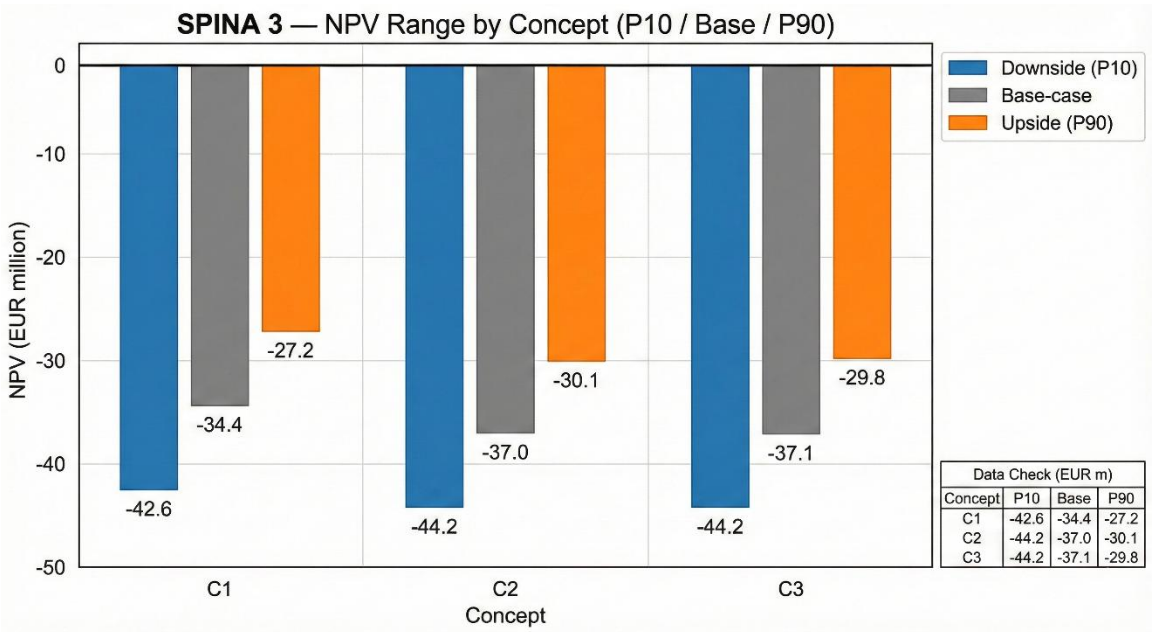

Figure 5.21 Scenario matrix visualization (3×3). Source: Author's Nano Banana Pro outputs, 2025.

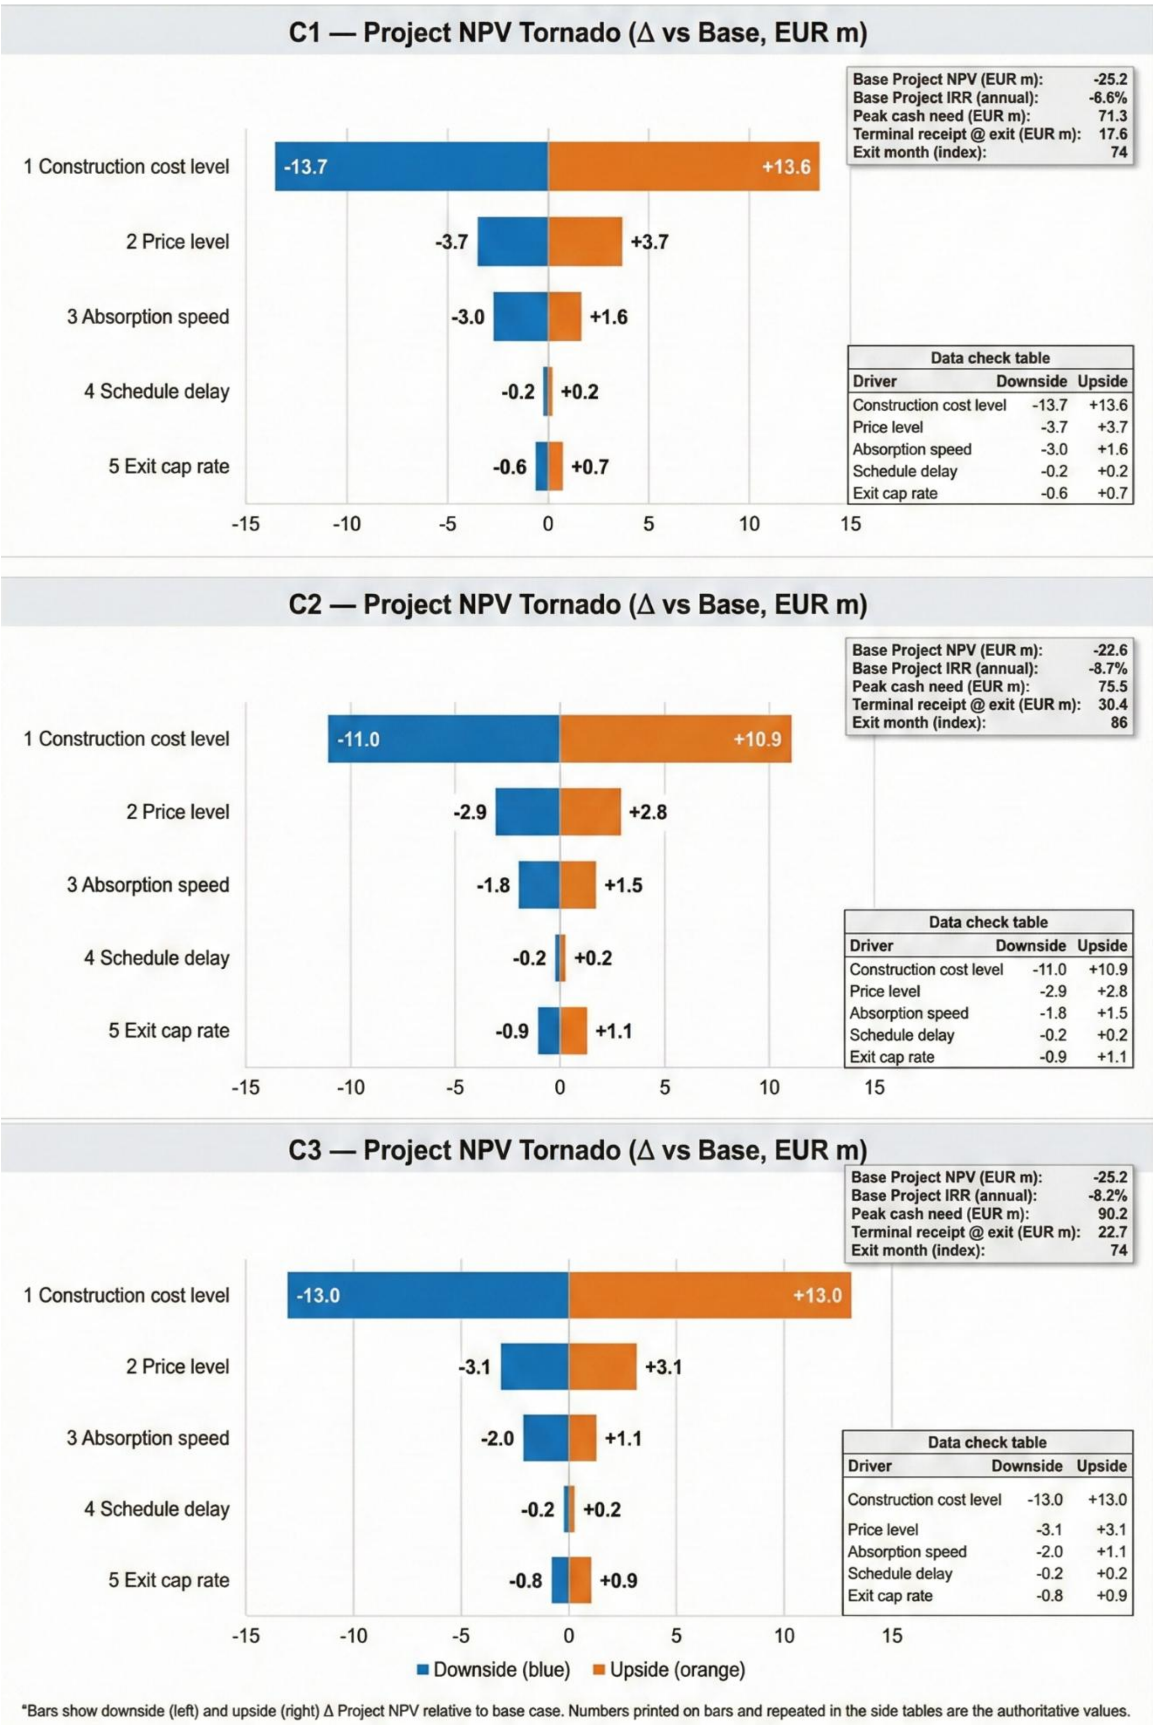

Figure 5.22 Tornado sensitivity charts. Source: Author's Nano Banana Pro outputs, 2025.

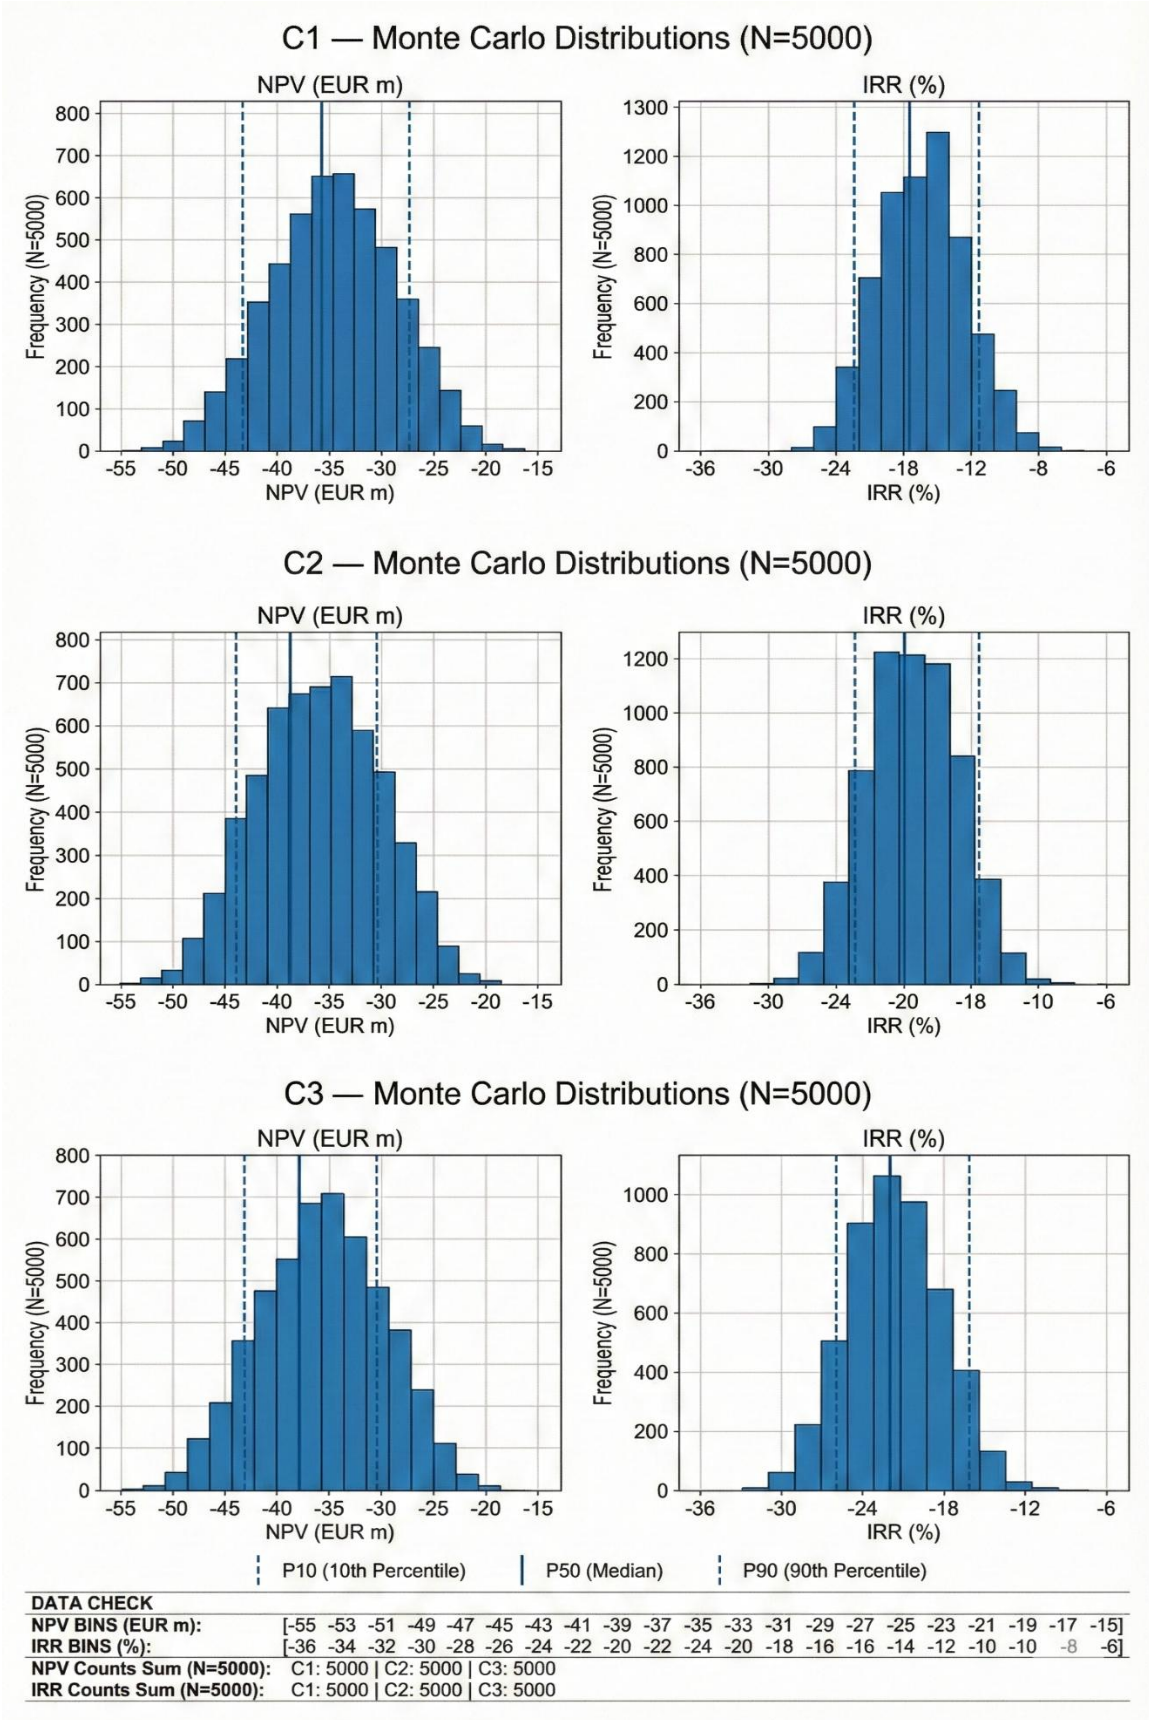

Figure 5.23 Monte Carlo distribution and CDF plots. Source: Author’s Nano Banana Pro outputs, 2025.

In architectural feasibility communication, decisions are often scrutinized through visual comparison as much as through tabulated KPIs. This post-processing step therefore supports transparency: by enforcing consistent scales and labelling, it reduces the risk of misinterpretation and makes the evidence base easier to audit.

5.3.6 Manual Adjustments and Final Delivery

This final step functions as a manual quality-assurance checkpoint that converts the iterated analytical outputs into a submission-ready chapter deliverable. In an architecture and built-environment thesis, the credibility of quantitative feasibility evidence depends on traceability (where each value comes from), interpretability (units and assumptions are unambiguous), and editorial consistency (figures, tables, and references allow the reader to follow the argument without confusion). Accordingly, the manual review focuses on verifying numerical consistency and constraint compliance, standardizing documentation conventions, and eliminating residual platform artefacts that can appear after multi-tool generation and export.

The manual review was applied to the complete decision-report package and its upstream analytical artefacts, with versioned filenames preserved to maintain the SSOT chain:

- Integrated Investment Decision Report – SPINA 3 Redevelopment (Turin, Italy) – Revised_v1.1.docx (Author's workflow artefact, 2025).
- Referenced analytical sources previously embedded in Sections 5.3.3–5.3.5 (financial model outputs, sensitivity charts, scenario matrix, risk register, and Monte Carlo results).

Quality-assurance checklist. The review applied a repeatable checklist targeting the most common failure modes in AI-assisted analytical writing and multi-file integration:

- Numerical consistency: repeated KPI values (NPV, IRR, peak cash need) match across summary tables, charts, and narrative statements.
- Units and magnitudes: currency units (EUR, EUR million) and area units (m^2) are stated consistently and avoid implicit conversions.
- Cross-references: all in-text references to Figure 5.X and Table 5.X match the final numbering sequence.
- Constraint consistency: reported concept metrics remain within the hard-constraint envelope described earlier in Chapter 5.
- Terminology: SSOT is used consistently; prohibited or legacy labels are removed; abbreviations are defined once and reused coherently.
- Figure/table captions: captions include a clear title and an explicit source note (file + page/sheet), and remain interpretable when read standalone.

5.4 How AI Decision Reports Support Spatial Iteration: Translating Feasibility Logic into a Phased Plan, Forma Massing, and Communication Outputs

This subsection demonstrates how an AI-generated real-estate decision report can actively guide spatial design iteration, provided that the workflow is disciplined and traceable. The report is not just used to "design buildings". It is also used to (1) extract decision-relevant constraints and risk messages, (2) lock a small set of core assumptions, and (3) translate them into a phased spatial plan that can be modelled in Autodesk Forma and communicated through a stylized master plan and context compositing.

5.4.1 Step A - Extract decision-relevant information from the AI report (not geometry)

The starting point is the Integrated Investment Decision Report – SPINA 3 – Revised_v1. The design workflow extracts only the information that can be converted into spatial decisions, for example:

- Which concept is financially safer and why (e.g., "residential-led" performs as the lowest-risk base).
- Which use-mixes are risky (e.g., large speculative office/commercial exposure is the highest risk).
- What the report recommends as risk mitigation (e.g., "phasing + go/no-go thresholds", avoid committing to uncertain non-residential too early).

What must remain aligned across iterations (single-source assumptions, auditability, stable phasing logic).This extraction produces an explicit "design translation brief": a list of hard constraints, risk warnings, and decision priorities.

5.4.2 Step B - Lock the anchor condition and the "base concept" (core assumptions)

From the extracted report logic, two core assumptions are locked and treated as non-negotiable inputs for design iteration:

- Base concept: adopt Concept 1 (residential-led) as the project "base plate", because it is the lowest-risk option across the three concepts (i.e., "least downside").
- Anchor condition: assume confirmed institutional cooperation (an anchor institution), which allows importing the most valuable part of Concept 3 (innovation campus / institutional anchor) without relying on speculative office demand.

In other words: Concept 1 provides the stable cashflow base; the institutional anchor converts part of Concept 3 into a de-risked third-sector program (teaching/R&D/incubation), rather than pure market office.

5.4.3 Step C - Translate the report into a "new planning scheme" that is spatially testable

Once the core assumptions are locked, the report conclusions are translated into a designable scheme with three layers: hard constraints → functional mix → spatial/typological rules.

Hard constraints (must align with the report)

Total buildable GFA ≈ 46,006 m², and the functional red lines must be respected:

Residential ≥ 40%; Office ≤ 40%; Retail ≤ 20%

The development must follow a phasing + risk gates logic: deliver the most certain components first; later-stage non-residential must be optional, reversible, and convertible.

The report warns against large speculative office/commercial (Concept 2 logic); the new plan must avoid creating a scenario where "vacant office blocks drag the whole project".

Recommended functional mix (under the institutional anchor condition)

The report-to-design translation adopts the following target mix (kept within the red lines):
Residential (including dormitory): ~58% (≈ 26,680 m² GFA)
Rationale: residential absorption is treated as the most stable demand base.
Institutional / innovation third-sector (teaching/R&D/incubation + convertible space): ~32% (≈ 14,720 m² GFA)
Rationale: use the anchor institution to replace speculative office risk with contracted demand.
Neighborhood retail: ~10% (≈ 4,600 m² GFA)
Rationale: limited to daily-life services and street activation; avoid oversized commercial exposure.

Typology rules that directly respond to feasibility risk

Dormitory (~3000 ㎡ GFA) is separated from the residence as an attached living function of the institution, which also responded to the increasing demand for student housing in the Turin market.
Institutional anchor building (≈ 9,000 m² GFA) is prioritized early to "make the cooperation real" and reduce uncertainty.
Incubator / co-working (≈ 3,500 m² GFA) is only expanded when pre-leasing signals exist.
A dedicated flexible building (≈ 2,220 m² GFA) is reserved as a risk buffer: it is designed so it can shift between training/small office → dorm/serviced apartments if the market weakens.
Retail is treated as shell space + phased fit-out: "sign first, build-out later", minimizing CAPEX and vacancy risk.

5.4.4 Step D - Convert the scheme into a phased delivery plan (Phase 0–3 with explicit scope and triggers) (Revised as requested)

The decision logic extracted from the AI report is implemented through a phased delivery plan. Here, phasing is not only a timeline, but a way to control feasibility risk: the project first delivers the most certain components, while later-stage non-residential programs are introduced only when clear demand signals and partnership commitments are in place.

Phase 0 (approvals and site preparation, 0–18 months)

Phase 0 focuses on enabling conditions and legal/technical prerequisites. The deliverables include: (i) the convenient master plan and the agreement for public facilities/works to be delivered, and (ii) confirmation of site constraints through environmental contamination checks and geotechnical investigation (the report identifies these as key prerequisites). This phase establishes the stable baseline against which all later phases can be evaluated.

Phase 1 (1–3 years): institutional anchor + first residential package + dormitory + main park skeleton

Phase 1 delivers the "certainty package" that operationalizes the institutional cooperation assumption and establishes early place identity. The scope includes:

- Residential: approximately 110 units, mainly 1BR/2BR, to prioritize absorbable product types.

- Dormitory: 150 beds, delivered in Phase 1 to create a stable daily population linked to the anchor institution.
- Institutional anchor building: approximately 9,000 m², delivered early to "make the cooperation real" in built form.
- Retail (starter package): approximately 1,000 m², oriented to daily services (e.g., convenience store + café type).
- Public space: the linear park main axis + a core square, delivered first at a "usable" level, with upgrades and finer finishes deferred to later stages.

Phase 2 (3–5 years): second residential package + incubator + supermarket and main-street retail

Phase 2 is only launched when the following conditions are met:

- Residential absorption (Phase 1): sales/lease performance reaches ≥70%.
- Anchor commitment: the institutional partner has signed a ≥10-year use agreement and/or a co-investment commitment.
- Incubator/co-working pre-commitment: documented demand (LOI/MOU) reaches ≥40% of the area.

With the triggers satisfied, Phase 2 expands both housing and partner-backed third-sector space, while strengthening daily-life services and public-space continuity. The scope includes:

- Residential: an additional 110 units delivered as the second housing package.
- Incubator / co-working: approximately 3,500 m², positioned as an extension of the anchor-led innovation program rather than speculative office.
- Retail: supermarket 1,600 m² + neighborhood retail/services 1,400 m², completing the main active street edges.
- Public space: completion of the second segment of the park, plus children's activity areas and slow-mobility connections to form a more continuous network.

Phase 3 (5–7 years): final residential package + flexible building + remaining retail and community functions

Phase 3 completes the neighborhood and embeds adaptability to respond to market uncertainty. The scope includes:

- Residential: 60 units, with the option to allocate larger unit types in this phase if market clarity and area values are stronger.
- Flexible building: approximately 2,220 m², delivered as an adaptability buffer (default use: training / small office; if demand weakens, conversion to dormitory or long-term rental housing remains feasible).
- Retail: remaining balance of approximately 600 m² to complete the local commercial offer.
- Public space: completion of the overall landscape system, including theme-based public art / science-education elements.

Across all phases, the principle remains consistent: each phase is defined by a clear scope of deliverables, and Phase 2 is protected by explicit start triggers to prevent premature exposure to uncertain non-residential demand.

5.4.5 Step E - Model in Autodesk Forma: keep massing aligned with the phased plan

Autodesk Forma is used to convert the phased plan into testable massing while keeping iteration disciplined:

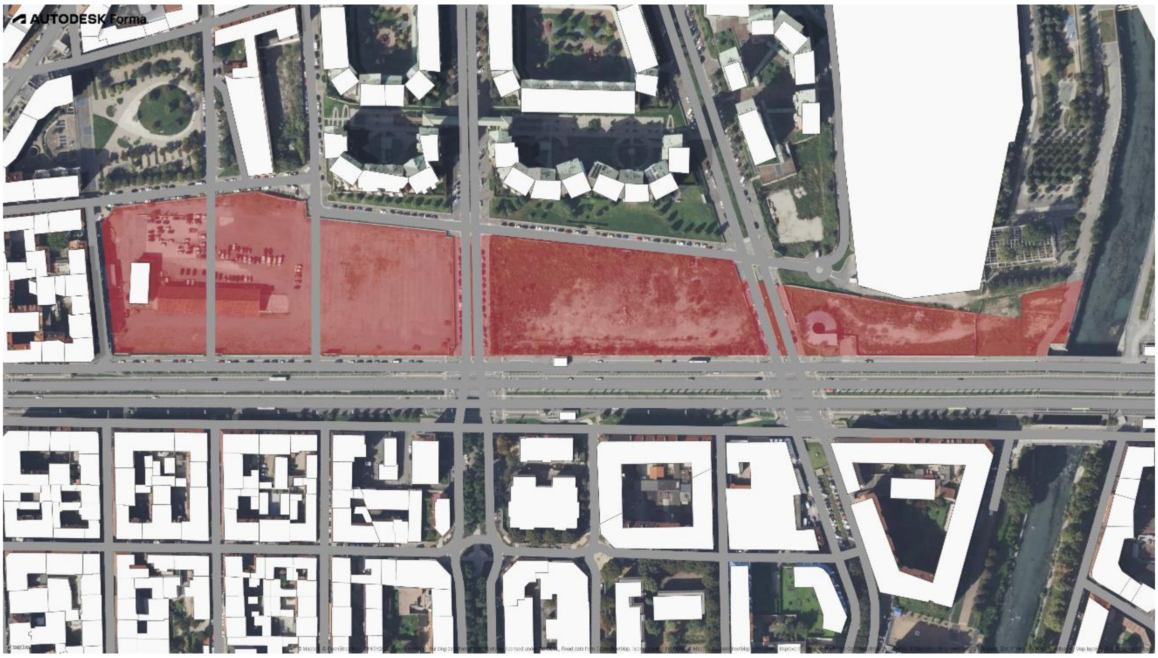

Figure 5.24 Phase 0 baseline (empty-site control map). Source: Author, exported from Autodesk Forma.

Figure 5.25 Phase 1 plan (residential + commercial + dormitory + institutional anchor + initial public space). Source: Author, exported from Autodesk Forma.

Figure 5.26 Phase 2 plan (residential + commercial + co-working/office + expanded public space). Source: Author, exported from Autodesk Forma.

Figure 5.27 Phase 3 plan (completion: low-density housing + flexible infill + public-space continuity). Source: Author, exported from Autodesk Forma.

Figure 5.28 Autodesk Forma bird's-eye output (massing + context). Source: Author, Autodesk Forma export.

5.4.6 Step F - Communication outputs: stylised master plan and context compositing (AI as visual assistant, not SSOT)

After massing is validated, Nano Banana Pro is used for communication steps, with strict controls:

Case 1: Competition-style stylized master plan

Goal: highlight the project boundary and reduce surrounding noise.
Critical rule: earlier colors (red/yellow/orange/blue) are functional coding only.

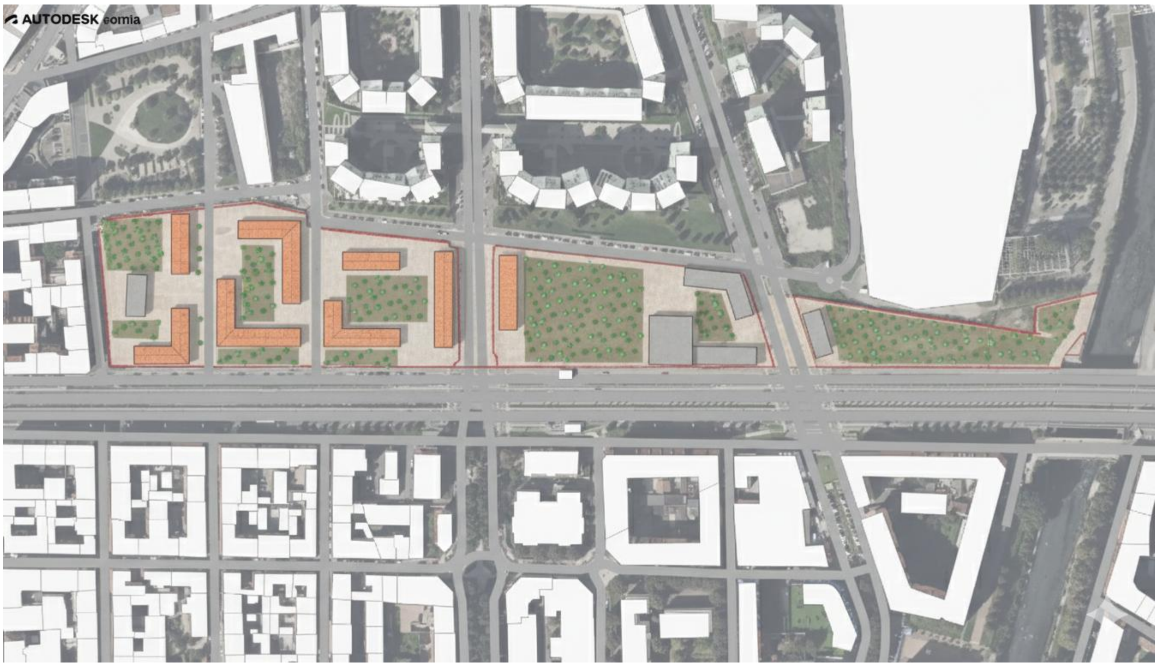

Figure 5.29 Stylized master plan (competition-oriented graphic; context de-emphasized). Source: Author, generated with Nano Banana Pro from Forma export.

Case 2: Context compositing into a realistic bird's-eye scene

Goal: test plausibility (scale, street continuity, shadows) and improve communication.
Controls: lock camera angle; mask edits to the site; preserve roads and surroundings; prevent global changes.

Figure 5.30 Context composite (proposal inserted into realistic bird's-eye scene). Source: Author, generated with Nano Banana Pro from base aerial + Forma massing.

Mini summary

This case shows three concrete benefits of using the AI decision report as an upstream "translation engine":

- A stable decision spine: Concept 1 provides a residential-led base; Concept 3 is imported only through an institutional anchor, avoiding speculative office risk.
- Risk-managed phasing: later packages are tied to explicit go/no-go triggers, and a conversion-ready building is reserved as a buffer.
- Faster, clearer communication: phased diagrams + Forma massing + stylized plan + context composite allows design intent to be explained consistently.
- At the same time, image-generation tools can introduce unwanted changes (blur, drifting boundaries, road deformation). Therefore, the workflow keeps a strict hierarchy:
- SSOT for geometry and quantities: phased plan + Forma model (author-controlled)
- AI outputs: communication artefacts only, always checked against the SSOT

In summary, the AI report supports spatial iteration only when it is used as a disciplined translator of feasibility logic into a testable phased design, rather than a direct generator of architectural form.

5.5 End-to-end Workflow Summary and Lessons Learned

This subsection brings together the empirical workflow described in Sections 5.3.1–5.3.6 and the spatial-iteration loop in Section 5.4 to provide a single end-to-end account. It explains, in plain terms, how the project moved from evidence collection to feasible concept comparison and finally to design communication: first by building a curated dataset, then by stabilizing key assumptions through a single source of truth (SSOT), and finally by iterating outputs with clear human audit checkpoints. The aim is to show how transparency and traceability were maintained throughout an architectural redevelopment feasibility-to-decision process.

In practice, the workflow operates like a controlled production line. It starts with an “input-side” stage where the project team curates a dataset that is specific to the site and the Italian regulatory/market context and then defines prompt logic to ensure that AI outputs remain tied to those inputs. Next, a single versioned SSOT document is used to lock the shared baseline—constraints, assumptions, and the definition of each concept—so that later iterations do not drift. Based on this fixed baseline, an “output-side” stage generates the downstream artefacts (tables, models, charts, and text). These outputs are not accepted immediately: at each handover between steps, a human audit checkpoint verifies basic correctness (units, internal consistency, and compliance with hard constraints) before the workflow proceeds.

This separation between input-side control and output-side production reduces unnecessary context noise, makes the prompt history easier to follow, and supports traceability through versioned deliverables, consistent with the governance principles established in Chapter 4.

Section 5.4 extends the same logic into spatial design iteration. Here, the AI real-estate decision report acts as a bridge between feasibility reasoning and design actions: it translates the financial/risk conclusions and development logic into concrete spatial priorities and testable moves (e.g., what to phase first, what to protect as non-negotiable, and where flexibility is allowed). The phased plan and the Autodesk Forma model then serve as the geometry-and-quantity SSOT—the authoritative reference for massing, areas, and spatial feasibility. Generative visual tools are used only for communication outputs (e.g., stylized plans and render-like graphics), and they must be checked against the phased plan/Forma SSOT to avoid visual “drift” being mistaken for verified geometry.

Table 5.15 End-to-end artefact map. Source: Author’s synthesis based on versioned workflow artefacts, 2025–2026.

Stage (section)	Main outputs	File names	What is 'locked' (SSOT)	What changes (iterations)
5.3.1 Concept Report	Single SSOT report consolidating constraints, market scan, and Concept 1-3 definitions.	Feasibility Study - Constraints, Market & Concepts (1. Cover & Version.docx)	Constraints, assumptions, and concept definitions used downstream as the narrative baseline.	Editorial clarifications and version increments until the SSOT is frozen for modelling.

Stage (section)	Main outputs	File names	What is 'locked' (SSOT)	What changes (iterations)
5.3.2 Modelling and Technical Feasibility Study	Autodesk Forma Proposals 1-3: concept-stage massing, area metrics export, and early environmental checks for comparative screening.	Autodesk Forma exports (Proposals 1-3) + embedded figures/tables in Section 5.3.2.	Concept-to-Proposal mapping and verified area metrics within the SSOT envelope (e.g., GFA cap and program proportions).	Geometry iterations to reconcile metrics and improve comparability; selection of representative outputs for documentation.
5.3.3 Financial modelling, sensitivity analysis, and risk assessment	Six-gate workflow (Gate 0-5) producing a comparable, auditable financial and risk evidence base for Concepts 1-3.	Gate 0-5 deliverables (see following rows).	Analysis boundary (pre-tax, project-level), shared modelling rules, and SSOT-aligned input discipline.	Only concept deltas (C1-C3) and declared uncertainty ranges; controlled re-runs with versioned patches.
Gate 0 - Setup and delivery framework lock-in	Directory structure, naming/versioning conventions, and binding analysis boundary definition.	Gate0_DeliveryDirectory_Versioning_AnalysisBoundary_PreTax_SPINA3.docx	Folder schema + file naming rules + scope boundary (pre-tax, project-level KPIs).	Minor wording clarifications before freeze; after freeze treated as governance policy (read-only).
Gate 1 - Model-ready parameters from SSOT	Locked hard constraints + concept program/phasing parameters + delta log defining allowed differences across C1-C3.	SPINA3_Inputs_HardConstraints_All; SPINA3_Inputs_ConceptProgramPhasing_C1-C3; SPINA3_Inputs_ConceptDeltaLog_C1-C3	Constraints register and concept definitions (quantities/phasing) used downstream.	Corrections only through versioned patches (no silent changes); deltas remain the only concept-varying inputs.
Gate 2 - Master Inputs and Data Gaps	Master Inputs table (deterministic baseline) + Data Gaps ranges/distributions; cost plan WBS; schedule baseline; cash-flow curve parameters; hold/operations assumptions.	SPINA3_Inputs_MasterInputs_All; SPINA3_Inputs_DataGapsRanges_All; Hold_Operations_AssumptionsPack; CostPlan_WBS; ScheduleBaseline_Gantt; CashFlowCurves_Params	Single consolidated assumption base + explicit uncertainty envelope; WBS and schedule baselines for comparability.	Range/distribution refinements and trace hotfixes; curve mappings updated only via controlled version increments.
Gate 3 - Base-case financial model runs	Same model structure executed three times (C1-C3) to produce base-case KPIs and embedded constraint/sanity checks.	SPINA3_FinancialModel_C1-C3_v1.0_20251227.xlsx	Model structure, KPI definitions, and constraint checks (same rules across concepts).	Re-runs only when upstream inputs are patched; outputs remain trace-linked to the frozen input set.
Gate 4 - Sensitivity and scenario testing	Tornado sensitivity outputs + 3x3 scenario matrix to identify dominant value drivers and downside exposure (same shocks across concepts).	SPINA3_Sensitivity_Tornado_v1.2.2_20251227_2pagesPerConcept_CleanFilled.pdf; SPA3_Scenario_Matrix_v1.0.1_20251227_Viz_v1.xlsx	Variable definitions and shock magnitudes; consistent presentation logic across concepts.	Visual layout refinements; scenario labelling/legend normalization; no changes to the underlying model outputs.

Stage (section)	Main outputs	File names	What is 'locked' (SSOT)	What changes (iterations)
Gate 5 - Risk register and Monte Carlo testing	Risk register (categorized, scored) + probabilistic simulations for C1-C3, reporting distributions and threshold probabilities.	SPINA3_RiskRegister_v1.0.4_20251227.xlsx; SPINA_MonteCarlo_Results_C1-C3_v1.0.2_20251227.xlsx	Risk taxonomy/scoring scheme and mapping between qualitative risks and quantifiable variables.	Simulation re-runs only when ranges/assumptions are patched; reporting improved through clarified captions and percentile notes.
5.3.4 Integrated comparative decision-making	Integrated decision report synthesizing constraints, technical feasibility, KPIs, sensitivity/scenario findings, and risk outputs into a conditional recommendation.	Integrated Investment Decision Report - SPINA 3 Redevelopment (Turin, Italy) - Revised_v1.1.docx	Evidence base remains tied to upstream artefacts (no new model layer or new numbers introduced).	Narrative tightening and audit-oriented clarification of conditions/mitigations; versioned human edits.
5.3.5 Supplement with visual charts	Chart standardization for comparability and communication (aligned axes, typography, and multi-concept layouts).	Nano Banana Pro exports (derived from the cited SSOT artefacts).	Underlying numeric values remain unchanged; only representation is adjusted.	Layout/typography refinement; consolidation of multi-concept visuals to reduce misinterpretation risk.
5.3.6 Manual adjustments and final delivery	Final quality-assurance pass: numbers, units, captions, cross-references, and reference list integrity; removal of export artefacts.	Human adjust Integrated Investment Decision Report - Revised_v1.pdf + final Chapter 5 DOCX package	Finalized figure/table numbering, terminology, and submission-ready formatting.	Correction of residual inconsistencies and layout issues; final polishing before submission.
5.4 Spatial iteration and communication loop	Phased delivery plan (Phase 0–3) translated from the AI decision report; Autodesk Forma massing aligned to the phased plan; communication outputs (stylized plan + context composite) treated as downstream representations, not evidence.	AI decision report (Section 5.3.4) + phased plan diagrams + Autodesk Forma phase exports + Nano Banana Pro visual composites.	Anchor condition and base-concept assumptions; phased plan scope per phase; Autodesk Forma model as geometry-and-quantity SSOT (author-controlled).	Visual styling and compositing only; any geometry change requires returning to the phased plan and Forma model (no silent drift in boundaries, roads, or footprints).

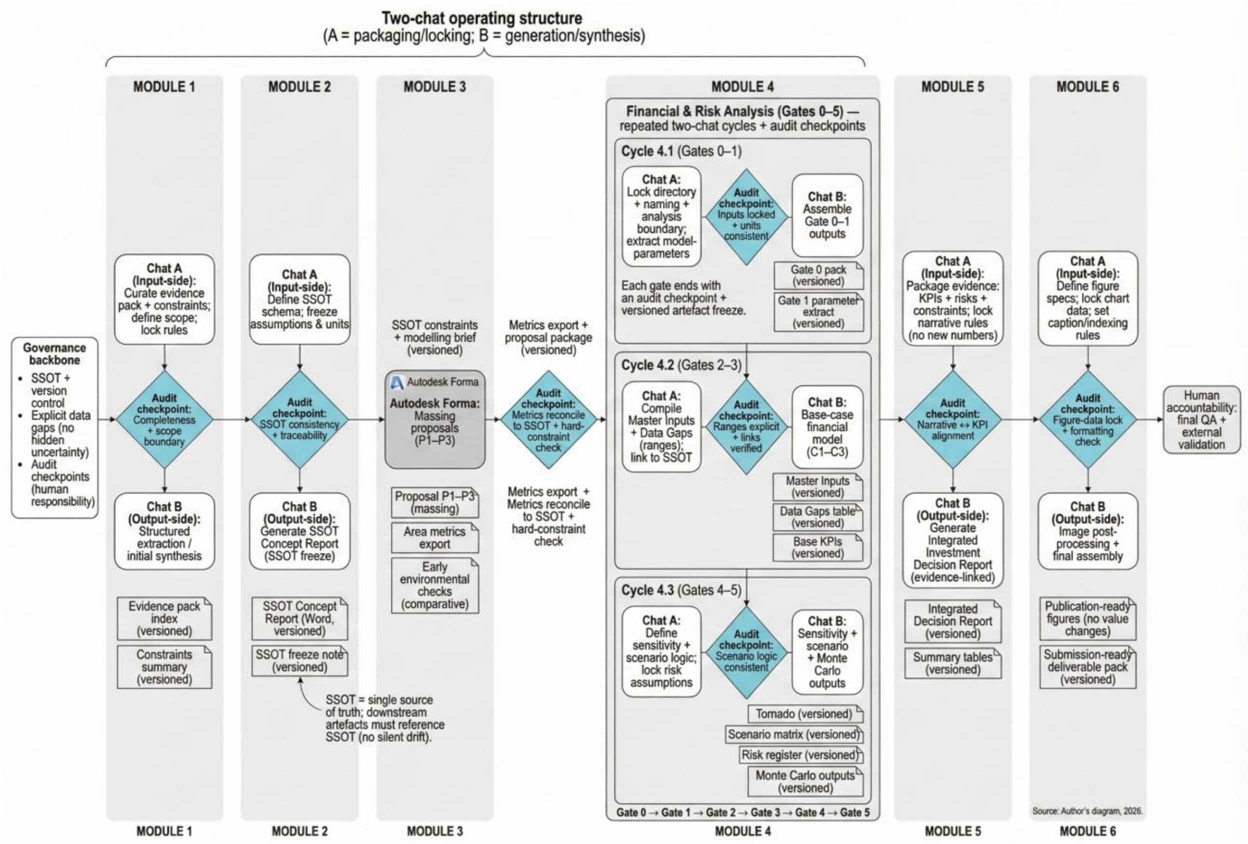

Figure 5.31 End-to-end empirical workflow (two-chat structure, audit checkpoints, and versioned artefacts). Source: Author's diagram, 2026.

Without introducing new results, the end-to-end map highlights the main evidence blocks generated in this chapter: (a) concept definitions and constraints consolidated in the Concept Report (Section 5.3.1); (b) concept-stage spatial and environmental validation produced in Autodesk Forma as Proposals 1–3 (Section 5.3.2); (c) a structured six-gate financial and risk workflow (Section 5.3.3, Gate 0–Gate 5), including base-case KPIs, sensitivity and scenario tests, and probabilistic risk outputs; (d) an integrated comparative decision narrative (Section 5.3.4); (e) publication-ready visual charts (Section 5.3.5); (f) manual quality assurance prior to final delivery (Section 5.3.6); and (g) a controlled spatial-iteration loop translating feasibility logic into a phased plan, Forma massing, and two classes of communication outputs (stylised plan and context composite), with explicit controls to prevent geometry drift (Section 5.4).

For built-environment research, the primary strength is traceability: each transformation from narrative concept to quantitative KPI is anchored to a versioned artefact with an explicit source note. Comparability is maintained by applying a consistent modelling structure across concepts, with differences confined to SSOT-defined deltas. Auditability is strengthened by explicit human checkpoints that address common failure modes in AI-assisted writing (mislabeled units, broken cross-references, and ambiguous sourcing), thereby improving the credibility of the resulting feasibility narrative. A further strength demonstrated in Section 5.4 is the extension of SSOT discipline to spatial truth: by treating the phased plan and Autodesk Forma model as the authoritative geometry-and-quantity layer, the workflow separates verifiable design states from downstream visual representations.

The workflow remains a concept-stage feasibility methodology: outcomes are dependent on market assumptions, cost plans, and operational parameters that carry uncertainty. The analyses do not claim regulatory compliance, detailed engineering coordination, or investment-grade valuation. Instead, they provide a structured basis for early decision-making and for identifying which assumptions require deeper verification in subsequent design development and due diligence stages. The spatial-iteration component (Section 5.4) is similarly schematic: Forma outputs support concept-stage plausibility checks, but they do not replace detailed planning review, design development, or stakeholder negotiation. Finally, multimodal and image-generation tools can introduce misleading drift (e.g., blurred boundaries or altered roads); this requires strict separation between SSOT geometry and communication layers and increases the operator skill requirement rather than eliminating it.

The approach is reusable for other urban regeneration and redevelopment feasibility cases where the decision problem requires both spatial validation and financial–risk reasoning. Transferable components include the SSOT-based input discipline, a reproducible directory and versioning convention, standardized KPI tables for concept comparison, and an explicit post-processing and manual QA step to ensure publication-ready outputs. Where spatial iteration is needed, Section 5.4 suggests a transferable pattern: use the AI report for extracting decision-relevant logic, translate it into a phased plan that can be model-tested, and assign a parametric modelling environment (e.g., Forma) as the geometry SSOT; treat image-generation as a downstream communication aid only, always audited against the locked spatial baseline.

In summary, Chapter 5 demonstrates that an AI-assisted workflow can be integrated into an architectural feasibility study in a controlled manner, if inputs are curated, outputs are versioned, and human verification remains an explicit part of the method. The addition of Section 5.4 clarifies that "end-to-end" integration must also include a controlled bridge from feasibility logic to spatially testable schemes: LLM outputs are most reliable as translation and structuring layers, while geometric truth must be anchored in an author-controlled model (phased plan + Autodesk Forma) and kept separate from downstream visual rendering. The next chapter builds on this empirical execution by reflecting methodological implications and broader applicability.

6. Discussion

This chapter discusses the empirical findings of Chapters 4 and 5 in relation to the research questions and the gaps identified in the literature. Rather than repeating the step-by-step workflow, it reflects on what changes when AI is introduced into early-stage planning and development feasibility: (i) how reliability, auditability, and accountability can be designed into an AI-assisted pipeline; (ii) what strengths and failure modes were observed in practice; and (iii) what implications follow for future research and professional adoption in an Italian/EU context; (iv)how AI-supported feasibility can be translated into design actions without losing traceability.

The empirical case study (SPINA 3 - Corso Principe Oddone, Turin) demonstrated that the central challenge is not the availability of individual AI features, but the governance of a multi-domain decision chain where spatial constraints, regulatory checks, market assumptions, and financial risk are tightly coupled. Accordingly, the contribution of this thesis is framed as a governed, end-to-end methodology: a large language model (LLM) is used primarily as an integration and documentation layer; parametric modelling (Autodesk Forma) anchors geometric feasibility and early environmental checks; and generative visual tools are treated as bounded communication aids. Crucially, Section 5.4 shows that the AI decision report can act as a disciplined translation layer between feasibility logic and spatial iteration, if geometry and quantities remain controlled by an explicit spatial SSOT and are checked at defined human audit points.

6.1 Critical reflection on the integration of AI into planning and development feasibility

From tool substitution to workflow governance. Chapters 2-3 show that many AI and digital tools in the built environment are optimized for narrow tasks. The case study confirms that this is a core limitation for feasibility work: the decision chain required planning interpretation, spatial reasoning, market positioning, cost and schedule logic, and financial risk analysis to be handled together. No single tool could do all of this in a way that was both usable and auditable. The practical response was therefore to design the integration as a workflow architecture (with controls, versioning, and reviews), rather than to search for one 'end-to-end' software product.

Why 'single-tool' support fails in practice. Two structural reasons explain why standalone tools struggled. First, real-estate decisions are multi-objective: a proposal must satisfy regulatory hard constraints, achieve spatial and functional coherence, and remain financially viable under uncertainty. Tools that perform well on one dimension (e.g., valuation dashboards or design automation) often cannot carry the assumptions required by the other dimensions without simplification or loss of meaning. Second, many vertical tools embed opaque models and proprietary data pipelines, which makes their outputs hard to audit, adapt, or defend - especially in an Italy-specific evidence environment and a student context. The methodology therefore prioritizes transparency over automation: each transformation from inputs to outputs must be visible and reviewable, even if this requires more user discipline.

In this sense, the case study reframes 'end-to-end' as a managed sequence of accountable states. The goal is not to automate every task, but to ensure that each decision-relevant state (constraints frozen, concept definitions fixed, model assumptions logged, risks quantified, and - in Section 5.4 - spatial consequences tested) is explicit, versioned, and capable of being challenged. This framing aligns with feasibility as a staged reduction of uncertainty: early outputs are conditional, and they gradually harden as new information becomes available and is verified.

Auditability as a feasibility requirement, not a reporting preference. In planning and development feasibility, auditability is not limited to citing sources in prose. It also includes traceable numerical assumptions, reproducible transformations across artefacts, and explicit scope boundaries that prevent silent drift. Chapters 4-5 operationalized this through a versioned SSOT (Single Source of Truth) and gate-based deliverables, combined with human audit checkpoints. This design aligns with information management principles that emphasize controlled information exchange and verifiable states of information (ISO 19650-1/2). It also responds to well-known spreadsheet risks, where small input errors or untracked changes can propagate into materially different outputs (Panko, 1998). The key reflection is that LLMs do not automatically create auditability; they can just as easily accelerate the production of plausible but untraceable narratives unless governance is enforced.

Why the two-chat structure matters. The workflow adopted a two-session structure (an input-side session and an output-side session) to reduce context noise and to manage platform constraints such as context-window limits and the tendency of long chats to accumulate irrelevant history. Conceptually, this separation proved useful because it forced a clean distinction between (i) configuring the evidence package and prompt logic, and (ii) generating and iterating outputs. When combined with SSOT freezing points, the two-chat pattern reduced assumption drift across iterations and made it easier to attribute each downstream table, model export, or chart to a specific upstream artefact version. In short, operational constraints became a methodological driver for stronger traceability and version control.

Gate-based controls as an operational translation of AI risk governance. The Gate 0-Gate 5 structure implemented in Chapter 5 can be read as a practical way to manage generative-system risks in a feasibility context. Gate 0 locks the delivery framework; Gate 1 extracts only the developable parameters required downstream; Gate 2 formalizes uncertainty through master inputs and explicit data-gap ranges; Gate 3 generates the base financial model under a fixed structure; Gate 4 stress-tests sensitivities and scenarios; and Gate 5 quantifies risk using probabilistic outputs and a risk register. Section 5.4 extends this logic downstream: once a decision report has been audited, it becomes an input to spatial iteration, but the 'truth' of geometry and quantities is held by the parametric model and its exports. This decomposition mirrors core ideas in the NIST AI RMF, where risks are managed through structured processes that map objectives and contexts, measure performance and uncertainty, and manage residual risk through controls and oversight (NIST, 2023).

Multimodality and the 'spatial truth' problem. Feasibility in the built environment is anchored in spatial facts: envelopes, access, daylight, setbacks, program fit, and other constraints that are often communicated through drawings, maps, and diagrams. The stress tests in Chapter 4 and the case-study work in Chapter 5 indicate that LLM performance is highly sensitive to the quality and structure of spatial inputs. When the model must infer geometry from text alone, failure modes include misreading constraints, proposing unrealistic spatial moves, and producing inconsistent cross-references. This limitation explains why Autodesk Forma is treated as a hard-constraint validation step, and why Section 5.4 positions the AI decision report as a bridge rather than a generator: the report can translate feasibility logic into design intentions and priorities, but the spatial iteration must be checked against a geometry-and-quantity SSOT (the parametric model, its phased massing, and exported KPIs). In other words, reliable feasibility requires a division of labor: LLMs are strong at structured reasoning and documentation, but they need an external constraint engine to anchor spatial claims.

Structured outputs as a second reliability bottleneck. The case study also suggests that end-to-end usefulness depends on the ability to create stable, machine-checkable artefacts (especially spreadsheets with formulas and PDFs with consistent figure/table indexing). Even when an LLM can describe a financial model correctly, feasibility requires that the model be reproducible, recalculable, and transferable across tools and reviewers. Where platforms do not reliably export formula-backed spreadsheets or consistent long-form documents in one run, the workflow must compensate through stricter SSOT definitions, intermediate artefacts, and controlled manual integration. This strengthens the argument that multimodality alone is not sufficient: concept-stage feasibility requires both spatial anchoring and robust structured-export capability.

Human agency and exogenous uncertainty remain decisive. Even when internal logic is well governed, feasibility depends on variables that sit outside the evidence package. Market information may be incomplete or time-sensitive; policy support and funding opportunities require communication with public agencies; and partnerships or anchor-tenant options depend on negotiation and institutional trust. LLM outputs can support these steps by preparing structured questions, summarizing positions, and highlighting sensitivities, but they cannot replace the real-world interactions that produce binding commitments. This boundary is why the workflow treats human review not as a 'manual fix', but as an accountability layer: a responsible actor must confirm which assumptions are verified, which remain ranges, and which decisions are provisional.

Accountability boundaries in feasibility recommendation-making. Recommendation quality depends on clearly defining decision rights. In professional feasibility and valuation contexts, credible recommendations require transparent assumptions, declared scope, and defensible reasoning, because stakeholders will challenge the basis of value and risk (IVSC, 2024; RICS, 2024). In the case study, the most defensible outputs were those that explicitly marked (i) what was sourced and verified, (ii) what was inferred from bounded ranges, and (iii) what required external confirmation. The discussion therefore supports a conservative stance: AI-assisted feasibility can strengthen documentation and analysis, but it does not shift responsibility away from the human decision-maker or the institution that adopts the workflow.

Governance and compliance in an EU context. The case study highlights that feasibility workflows increasingly intersect with data protection, confidentiality, and organizational governance. Uploading project documents to external platforms raises questions of lawful basis, purpose limitation, and data minimization under the GDPR (European Union, 2016). In parallel, the EU AI Act formalizes risk-based requirements for many AI systems, reinforcing expectations around transparency, oversight, and documentation (European Union, 2024). These constraints strengthen the value of the evidence-package discipline proposed in Chapter 4: limiting uploads to what is necessary, maintaining a clear data inventory, and separating sensitive inputs from generated outputs. The broader implication is that 'AI integration' is not only a technical optimization; it is also an organizational and legal design problem that shapes what can be automated and what must remain human-controlled.

6.2 Strengths and weaknesses observed

Strength 1 - Traceability and comparability through versioned artefacts. The strongest empirical benefit of the methodology is the ability to keep a complex, multi-step process auditable. Each transition from narrative concept to quantitative KPI was anchored to a versioned artefact (SSOT, master inputs, data-gap ranges, financial model exports, sensitivity/scenario outputs, and risk registers). Comparability across the three options was maintained by applying a consistent modelling structure and limiting differences to SSOT-defined deltas. This approach reduced the 'moving target' problem common in early-stage feasibility, where different options are often evaluated with subtly different assumptions and therefore cannot be fairly compared.

Strength 2 - Explicit error-control against hallucination and drift. The workflow did not assume that model self-consistency equates to correctness. Instead, it treated hallucination and drift as expected risks in generative systems (Bender et al., 2021; Ji et al., 2023) and implemented controls: freezing the SSOT, using gated deliverables, and inserting human audit checkpoints that check units, cross-references, and constraint compliance. In practice, these checkpoints reduced downstream rework because errors were detected at step boundaries rather than after a full report was compiled.

Strength 3 - A productive division of labor across tools, including a feasibility-to-design bridge. A key benefit was the clarity of role assignment within the toolchain. The LLM was used where it adds unique value: synthesis of heterogeneous evidence, structured narrative reasoning, generation of comparable tables, and orchestration of the gate sequence. Section 5.4 extends this by using the audited AI decision report as a translation layer: it turns feasibility reasoning into concrete spatial priorities (phasing logic, program placement, risk-aware design choices) that can then be tested in a parametric model. Autodesk Forma was used where text-based reasoning is weakest: geometric validation and early environmental checks, treated as the geometry-and-quantity SSOT for the iterative loop. Generative image tools were used only at the end to improve the legibility of charts and diagrams, without changing underlying data; this visual step required explicit data locking to avoid persuasive but inaccurate representations. This division of labor avoids the common failure mode in which one tool is pushed beyond its competence boundary.

Strength 4 - Reduced coordination friction for small teams. Although the workflow does not eliminate professional responsibilities, it can reduce coordination friction in a small-team or student context. By consolidating assumptions and outputs into a coherent artefact chain, it becomes easier to communicate what has been decided, what remains uncertain, and what evidence supports each conclusion. This reduces unnecessary iteration caused by inconsistent briefs, missing inputs, or undocumented changes between stakeholders.

Strength 5 - Better preparation for external decision conversations. Gate outputs can be repurposed as briefing documents for interactions that AI cannot perform directly. For example, sensitivity and scenario outputs help define which market variables or policy conditions would most change feasibility, informing what evidence to prioritize and what questions to ask in stakeholder meetings. Similarly, the risk register provides a structured basis for negotiating responsibilities and mitigation measures with external parties. In this way, the workflow increases the quality of human decision-making not only through analysis, but also through better preparation for negotiation and institutional coordination.

Weakness 1 - Dependence on multimodal capability and structured exports. The workflow quality remained constrained by the platform's ability to interpret multimodal evidence and to output stable, machine-checkable artefacts. In concept-stage feasibility, reading spatial evidence (plans, maps, regulatory diagrams) and producing recalculable spreadsheets or consistent PDFs is not a convenience; it is a condition for end-to-end reliability. Where native exports were limited, the methodology had to compensate through modular outputs and manual integration. This limitation also applies to communication imagery: generative visuals can improve readability, but they can also introduce subtle distortions unless they are explicitly checked against locked data and geometry exports.

Weakness 2 - Cost and compute budgeting as hidden time costs. A practical limitation observed during the empirical work is that long, complex sessions can become inefficient: compute requirements increase with conversation length, interface performance can degrade, and complex artefact generation can fail or require repeated runs. This creates an operational trade-off: while AI accelerates many micro-tasks, the workflow must be carefully segmented to avoid losing time to platform latency, timeouts, or resource exhaustion. The methodological implication is that switching sessions at key milestones is not merely convenient; it is part of compute budgeting and therefore part of feasibility scheduling.

Weakness 3 - The operator skill requirement increases, not decreases. The method reduces drafting workload, but it increases the need for structured prompting, domain judgement, and careful review. The operator must understand planning constraints, market logic, and financial modelling well enough to detect implausible outputs, enforce scope boundaries, and interpret sensitivity and risk results. In other words, the workflow can compress some labor into fewer roles, but it raises the competence and attention required of those roles.

Weakness 4 - External validity is limited by the evidence boundary. The case study is concept-stage and therefore necessarily conditional. The most credible outputs were those anchored to verifiable evidence in the uploaded package and those validated through parametric modelling. Where information was missing (for example, precise policy commitments, partner terms, or transaction-specific market comparables), the workflow relied on explicit ranges and scenario structures rather than false precision. This improves transparency but does not remove uncertainty. The outputs should therefore be understood as decision support for early-stage option screening, not as investment-grade underwriting.

Weakness 5 - Platform dependency and model drift. The stress testing in Chapter 4 already noted that reproducibility can be constrained by model updates, shifting platform features, and differences between access tiers. In practice, this means a workflow that is stable today may require re-validation tomorrow, and prompts may need to be recalibrated when models change behavior. The two-chat architecture and SSOT versioning mitigate this risk by making each run explicit, but they do not eliminate it. This highlights a broader methodological point: AI-assisted feasibility requires not only artefact versioning, but also periodic workflow re-testing as part of quality assurance.

Weakness 6 - Privacy and confidentiality constraints on evidence packaging. Feasibility work often involves sensitive information (ownership data, transaction terms, tenant discussions, proprietary costs). Even in an academic project, data minimization and purpose limitation constrain what can be uploaded and therefore what the model can 'know'. This introduces a tension: increasing evidence completeness can improve analytical quality, but it can also increase compliance risk. The workflow partially resolves this by separating what is essential for reasoning (constraints, ranges, anonymized assumptions) from what is sensitive and should remain external to the AI platform (European Union, 2016; European Union, 2024).

6.3 Methodological implications

A minimum viable governance stack for AI-assisted feasibility. The empirical results suggest that the critical innovation is not a particular prompt, but a small set of governance components that make AI outputs usable in a decision context. These components include: (1) an evidence package with a data inventory; (2) a versioned SSOT with explicit freezing points; (3) gate-based deliverables with acceptance criteria; (4) an audit checklist for numerical consistency and referencing integrity; (5) a defined responsibility boundary that clarifies what AI outputs can and cannot justify; and, where spatial claims matter, (6) an explicit spatial SSOT (a parametric model and exported KPIs) against which narrative recommendations must be checked. Together, these elements convert generative output into a controlled artefact chain that supports review, comparison, and later design development.

SSOT design as a schema, not a document. One methodological lesson from the case study is that SSOT works best when treated like a small database schema rather than a narrative report. Inputs should be typed (units, currencies, dates, and scope), tagged by source, and linked to downstream artefacts. Where the workflow spans multiple sessions, versions should be incremented only at defined freeze points, and patches should be recorded as deltas rather than silent edits. With the addition of Section 5.4, this schema logic also applies to spatial data: the parametric model and its exports become the authoritative reference for quantities and early environmental KPIs, while the AI decision report remains a controlled narrative layer that must stay consistent with those outputs.

Implications for practice: designing handoffs between narrative and models. In professional workflows, feasibility studies often fail not because a model cannot be built, but because assumptions and constraints are not consistently transmitted between disciplines. The SSOT and gate structure provide a practical interface for handoffs: they specify which variables are fixed, which are ranges, and which are pending external confirmation, and they preserve a record of when each assumption changed. By treating the audited decision report as a controlled brief for modelling, and by treating the parametric model as the spatial SSOT for checking, the workflow reduces the risk that design iteration quietly departs from the feasibility basis. This approach is compatible with established valuation and risk management standards, which emphasize transparent assumptions and clear scope definitions (IVSC, 2024; RICS, 2024). However, adoption would still require organizational policies on data governance and review responsibilities, especially when external AI platforms are used.

Implications for project management: feasibility as a scheduled sequence of freeze points. The workflow suggests a different way to plan feasibility work. Instead of treating analysis as a single report-writing activity, the project can be managed as a sequence of acceptance milestones: evidence package completion, SSOT freeze, parametric validation, base-case financial model, sensitivity and scenario set, risk quantification, and - where relevant - a spatial-iteration loop with explicit check points. Each milestone produces artefacts that can be reviewed by decision-makers, and each establishes the conditions for the next step. This milestone logic makes the process more resilient to platform interruptions and staff changes, because the project state is represented by artefacts rather than by tacit knowledge.

Implications for education: shifting skills from drafting to auditing. For architecture and real-estate education, the method suggests a shift in emphasis. Students should not be trained only to 'use AI tools', but to define decision problems, construct evidence packages, design SSOT structures, and audit outputs for traceability and plausibility. This includes critical literacy about AI-generated images and diagrams: visuals can strongly shape judgement, so students must learn to verify charts, figures, and spatial claims against the underlying data and geometry exports. This competence is close to professional practice: it mirrors how feasibility work is defended in front of clients, public authorities, or investment committees, where credibility depends on explainable assumptions and documented constraints rather than on narrative fluency.

Future research directions. Five directions are prioritized by the case study. First, the development of more reliable multimodal workflows that can interpret and cross-check spatial evidence (plans, maps, regulatory diagrams) without forcing manual translation into text. Second, more robust structured-export pipelines (formula-backed spreadsheets, auditable PDFs, and stable long-form report generation) that reduce the integration burden for complex models and documents. Third, reproducibility studies that measure output stability under platform updates and model drift, using repeated-run protocols and explicit error taxonomies. Fourth, research on governance mechanisms for AI-generated visuals and diagrams, where the risk is not only factual error but also persuasive misrepresentation. Fifth, governance research on how responsibility boundaries should be formalized when AI is used for decision support in development feasibility, particularly under EU regulatory expectations. These directions can be tested using controlled versions of the evidence package and gate workflow, and they align with risk-based governance frameworks for AI (NIST, 2023; European Union, 2024).

6.4 Forward-looking implications for professional practice and the real-estate sector (2025-2026 horizon)

From breakthrough to deployment: why the timing matters. The year 2025 marks a visible shift from generative AI as a novelty to generative AI as a workplace infrastructure. Survey evidence suggests that AI use is now widespread but uneven: many organizations report regular AI use, yet only a smaller share report scaling it deeply into core workflows (McKinsey & Company, 2025). In parallel, broader indicators show rapid diffusion and growing investment momentum (Stanford University Institute for Human-Centered Artificial Intelligence, 2025), and labor-market tracking suggests that a majority of working-age adults have now used generative AI in some form, with work use rising more slowly than nonwork use (Bick et al., 2025). These signals matter for the built environment because feasibility and design coordination sit in the middle of multiple professional services, where partial adoption can create new coordination gaps: one actor may draft and iterate faster with AI, while another operates on traditional cycles and expects conventional documentation.

Potential benefits for feasibility and redevelopment practice. If adopted with appropriate controls, AI can reduce friction in evidence-heavy work that dominates early redevelopment phases: consolidating heterogeneous regulations and site evidence into a brief, maintaining consistent assumption sets across options, generating comparable scenario narratives, and preparing stakeholder-facing documents. In this thesis, the most credible gains came from treating AI as an integration layer that produces checkable artefacts, rather than as a substitute for modelling. In professional settings, this points to a near-term value proposition in 'documentation velocity' and 'option screening': faster iteration cycles, clearer traceability of assumptions, and better preparation for negotiations (tenants, funders, municipalities) - especially when the outputs remain anchored to SSOT discipline and parametric checks.

Risks and negative externalities. The same 2025 diffusion trend also amplifies familiar risks. First, adoption is likely to be uneven within firms and across the supply chain; evidence from enterprise usage suggests a widening gap between 'frontier' and median users, especially for advanced capabilities (OpenAI, 2025). In real-estate decision-making, such gaps can translate into asymmetric control over narratives and assumptions, where those who operate faster also shape what is considered 'reasonable'. Second, AI-generated outputs can increase the volume and polish of documentation without increasing its truthfulness, raising the risk of overconfidence and persuasion-by-format. Third, model drift and platform dependency mean that governance must be treated as an ongoing quality-assurance practice, not a one-off setup. Finally, privacy and IP constraints remain material: feasibility work often involves sensitive transaction terms and third-party data, which may limit what can be uploaded and what can be automated under GDPR and contractual obligations.

How to interpret 'AGI in 2026' discourse. Several leading AI-company figures argue that capabilities may continue to accelerate rapidly, with some explicitly discussing 2026 as a plausible horizon for systems that can generate novel insights or 'powerful AI' (Altman, 2025; Amodei, 2024). These timelines are speculative and should not be treated as forecasts for the built environment. However, they are useful as scenario prompts: if capabilities jump, the critical bottleneck in redevelopment will still include physical constraints, regulatory approval pathways, and responsibility for decisions that have legal and financial consequences. In other words, even very capable models do not remove the need for a geometric and regulatory ground truth, nor do they resolve accountability. The governance design tested in this thesis - evidence packaging, SSOT freezing, gated deliverables, and explicit human decision rights - becomes more important, not less, as model capabilities (and persuasive power) increase.

A constructive stance for architects and developers. A forward-looking but critical position is therefore to treat AI as an organizational capability that must be engineered, audited, and trained - similar to how BIM and information management matured through standards and defined roles. Near-term professional adoption should focus on workflows that (i) keep assumptions explicit and versioned, (ii) use parametric or BIM environments as spatial SSOTs for quantity and environmental checks, (iii) restrict generative imagery to post-processing that is checked against locked data, and (iv) define clear responsibility boundaries for recommendations. Within the EU context, this stance is also pragmatic: it anticipates tighter expectations around traceability, data governance, and oversight as AI tools become normalized in professional practice.

Overall, the discussion supports a pragmatic interpretation of 'end-to-end' AI integration. Feasibility can be accelerated and made more transparent when AI is treated as a governed integration layer that produces auditable artefacts, and when spatial claims are anchored in a parametric or BIM-based SSOT. The addition of the Section 5.4 loop reinforces this point: an audited decision report can help bridge feasibility logic to spatial iteration, but only if geometry and quantities remain verifiable and human-audited. Looking ahead, wider AI adoption in 2025-2026 is likely to increase both the opportunity (faster evidence handling and option screening) and the risk (persuasive but unreliable outputs, uneven adoption, and governance debt). Credibility and accountability therefore remain grounded in explicit controls and in the continued role of human expertise in verifying constraints, managing uncertainty, and obtaining external commitments.

7. Conclusions

This thesis examined how AI can be integrated into early-stage real-estate decision-making within an Architecture program. Using the SPINA 3 – Corso Principe Oddone case study in Turin, the research tested a practical workflow that combines large language models (LLMs) with governance controls, parametric validation, and human review across a full early-stage feasibility and decision-support cycle.

A central claim of the thesis is that valuation, when supported by AI and governed properly, can act as a bridge between spatial design and feasibility. In this view, economic evaluation is not an ex-post technical check that arrives after spatial design and technical feasibility analysis. Instead, it is a design-support process that runs alongside concept development. It helps translate design choices—program mix, massing, phasing, delivery strategy—into comparable consequences such as cost, timing, risk exposure, and decision indicators, and then feeds those consequences back into the next design iteration. This bridging role is especially important in architectural decision-making because spatial quality, constraints, and viability must be considered together, not in sequence.

The empirical workflow shows that this bridge only works when the information chain remains transparent and stable. For this reason, the methodology relies on a curated evidence package, a Single Source of Truth (SSOT) with explicit freeze points, staged Gate 0–5 deliverables, audit checkpoints, and version control. Autodesk Forma is introduced as a parametric validation step to reduce spatial errors, and to ground key economic inputs in measurable geometry, and can act as a preliminary validation platform to facilitate the iteration of spatial design. Generative image models are used only downstream to improve communication-quality figures, without altering the SSOT or changing numeric outputs.

7.1 Answers to research questions

RQ1: To what extent can AI support an end-to-end feasibility and investment decision process for a complex urban redevelopment project in Italy?

The thesis concludes that AI can meaningfully support an end-to-end process when "end-to-end" is defined as a sequence of controlled, auditable deliverables rather than a single autonomous run. Within the evidence boundary established by the curated dataset and SSOT, the AI system can draft structured outputs, compute consistent comparisons across alternative concepts, and assemble a coherent decision package. In the SPINA 3 case, this made it possible to move from constraints and concept assumptions to comparable feasibility outputs and risk-informed interpretations with less rework and clearer traceability.

At the same time, the workflow is not autonomous in any real-world sense. Key assumptions that depend on external reality, policy interpretation, financing terms, partner commitments, and market commitments, still require human responsibility and, in many cases, external stakeholder validation.

The method therefore positions AI as a decision-support accelerator, not a substitute for professional judgement, institutional negotiation, or accountability.

RQ2: What governance mechanisms are required to mitigate hallucination, drift, and inconsistency, while preserving auditability and accountability?

The thesis shows that governance mechanisms are not optional add-ons; they are the enabling conditions that make AI outputs usable in appraisal and design support. In practice, the most effective controls were SSOT freezing to prevent silent changes in assumptions, staged Gate deliverables that create clear checkpoints and acceptance criteria, audit checkpoints that catch unit errors and scope drift, and strict version control that supports patching without destabilizing the workflow.

The two-chat operating structure also proved important. By separating an input-side session - where rules, evidence boundaries, and packaging are locked, from an output-side session, where drafting and assembly occur, the method reduces context overload and helps maintain consistent logic across long tasks and platform constraints. Together, these controls allow valuation results to remain stable enough to inform design iterations instead of becoming a moving target.

RQ3: What are the observed strengths, limitations, and role boundaries of an LLM-centered workflow in the real-estate feasibility domain?

Empirically, the strengths of the approach lie in speed, structure, and synthesis. LLMs are effective at turning heterogeneous inputs into well-organized artefacts, identifying gaps, maintaining consistent comparative narratives across concepts, and supporting disciplined reasoning when the evidence boundary is explicit. This can improve productivity and clarity, especially in early-stage feasibility where uncertainty must be stated rather than hidden.

However, limitations remain significant. Without strong controls, hallucination and drift can reappear, particularly in long periods, under weak evidence, or when outputs require complex structured exports such as formula-correct spreadsheets or perfectly consistent cross-references. Spatial understanding also remains fragile if geometry is treated only as text. For these reasons, the workflow requires parametric validation and human auditing, and it must treat feasibility conclusions as conditional decisions rather than deterministic claims.

7.2 Key findings

Across Chapters 2–6, the study supports a clear finding: single-tool "end-to-end" AI support is not realistic for real-estate decision-making in a complex urban redevelopment context. The challenge is not only technical capability, but also responsibility. Feasibility work depends on evidence, scope boundaries, and stakeholder commitments that cannot be safely delegated to an opaque system.

The practical contribution of the thesis is therefore methodological rather than tool-specific. The workflow demonstrates how an LLM can be used as an integration layer, linking appraisal logic, documentation, and comparative reasoning, while governance controls keep outputs auditable. In this configuration, valuation becomes a design-support bridge. It accompanies spatial decisions by translating program, massing, and phasing choices into comparable economic and risk implications, which can then inform further spatial refinement.

The case study also shows why parametric validation matters in an Architecture context. By using Autodesk Forma to generate measurable massing options and exports, the workflow reduces the risk that feasibility inputs are detached from spatial reality. This improves the quality of the appraisal conversation: design options are not judged only by narrative plausibility, but also by quantified geometry that can be reconciled to the SSOT and fed consistently into downstream gates.

Finally, the study confirms that communication tools should be treated as downstream layers. Generative image models can improve the clarity and presentation of figures, but they must not be allowed to change values, assumptions, or conclusions. In an auditable appraisal workflow, visuals are outputs of a locked dataset, not a substitute for it.

7.3 Limitations

Several limitations should be recognized. First, the empirical evidence is based on a single case study and a student-driven workflow. While the governance principles are transferable, the specific market assumptions, policy constraints, and stakeholder environment are context-dependent.

Second, the method is bounded by the available evidence. When key inputs are uncertain or require real-world agreements, the workflow can only represent them as explicit ranges and conditional scenarios. This is useful for decision-making under uncertainty, but it cannot replace external verification.

Third, platform constraints affect reproducibility. Long conversations can degrade performance, and structured export reliability remains uneven. This increases the importance of staged artefacts, freeze points, and patching, but it also adds overhead.

Finally, the approach still relies on operator competence. The methodology reduces manual workload, but it raises expectations for audit discipline, basic financial literacy, and the ability to recognize when the model is stepping beyond the evidence boundary.

7.4 Looking ahead: AI acceleration and the role of architectural practitioners

The AI landscape is developing rapidly, and public debates increasingly include claims that very general systems, sometimes labelled “AGI”, could arrive soon, even within 2026. The thesis does not depend on any single prediction, and it is not possible to state a reliable timeline. What can be stated with confidence is the direction of travel: AI systems are becoming more capable, more multimodal, and more integrated with other tools. This will change how feasibility work is produced, and it will also increase expectations for control and responsibility.

This raises a common question in the built-environment disciplines: will architects and related professionals be replaced? The case study suggests a more precise answer. Many tasks are likely to be automated or heavily assisted, such as drafting documents, extracting and structuring inputs, generating option comparisons, producing sensitivity outputs, and preparing report-ready visuals. These are meaningful changes that can reduce routine work and shorten cycles.

However, replacement is not the most helpful frame. Real-estate decision-making is shaped by responsibility, negotiation, and context. Even if models become stronger, the work still requires someone to define the scope, validate units and constraints, accept or reject assumptions, and take accountability for decisions. It also requires engagement with external stakeholders. Policy support, financing conditions, and partnerships are obtained through institutional processes and human relationships, not only through computation.

In this sense, AI shifts the practitioner's role rather than removing it. The architectural professional becomes more of an appraisal-literate integrator: someone who can connect spatial choices to feasibility consequences, curate evidence responsibly, supervise governance controls, and use appraisal outputs to guide design and phasing decisions. The methodology proposed in this thesis can be read as preparation for that future role. It turns AI from a "black box assistant" into a governed workflow where design support, traceability, and accountability remain visible.

7.5 Recommendations for practice and future research

For practice, the main recommendation is to adopt a minimum governance stack as standard operating procedure when AI is used in appraisal. This includes maintaining a curated evidence package, defining an SSOT with freeze points, and working through staged gates with acceptance criteria and audit checkpoints. Parametric modelling should be used as a validation step for geometry and constraints, while generative images should be treated strictly as a communication layer based on locked data. Most importantly, valuation should be used as a design-support process. Appraisal checkpoints should be scheduled alongside design milestones so that feasibility evidence can inform spatial decisions in time, rather than arriving only after design has already solidified.

For future research, two directions are especially relevant. First, the empirical workflow could be extended to document, in more detail, how AI-supported decision outputs feed back into spatial design and phasing plans, and how those design revisions are then re-tested through the gated appraisal loop. Second, more work is needed on reliable structured outputs—such as formula-correct spreadsheets, consistent cross-references, and robust multimodal handling—together with privacy-aware evidence handling suitable for real projects.

In conclusion, the thesis argues that AI can strengthen project appraisal in architecture when it is governed as a traceable method rather than adopted as a single tool. Under this condition, valuation becomes an active bridge between design and feasibility. It supports spatial decision-making by making trade-offs explicit, uncertainty visible, and responsibility clear.

Attachments

ID	File / Artefact	Format	Role in thesis	Primary use (chapter/section)
A1	**SPINA 3 CORSO ODDONE** (Concept / constraints report)	DOCX	Case-study baseline: site context, constraints, and early concept assumptions.	Ch.5 (Case study inputs)
A2	**Cover & Version** (concept definitions and program metrics; SSOT-like)	DOCX	SSOT-like concept and program definition used for downstream gates and cross-file consistency.	Ch.5 Gate 1–2 (SSOT extraction + parameterization)
A3	**Environmental & infrastructure visual data**	PDF	Site context evidence and spatial constraints (transport, environment) supporting the feasibility narrative and spatial checks.	Ch.5.1–5.3 (context + constraints evidence)
A4	**OMI price & rent dataset** (selected zones, 2016–2025)	XLSX	Local market evidence for price/rent assumptions and ranges used in market module and financial inputs.	Ch.5 (market assumptions + ranges)
A5	**Gross Domestic Product** (Torino / relevant territorial unit)	XLSX	Socio-economic context: macro demand-side plausibility check and triangulation for market narrative.	Ch.5.1–5.2 (context + plausibility checks)
A6	**Employment rate & education level** (Torino)	XLSX	Socio-economic context: labor/education profile supporting demand logic and target segmentation plausibility.	Ch.5.1–5.2
A7	**Population within the City of Torino**	XLSX	Demographic baseline to support demand-side reasoning and scenario plausibility.	Ch.5.1–5.2
A8	**Surrounding area data 2021 & Torino municipal projections**	XLSX	Demographic / territorial projections to support trend narrative and uncertainty framing (ranges).	Ch.5.1–5.2
A9	**Household size & disposable income**	XLSX	Demand-side affordability / household structure context; supports	Ch.5.1–5.2

			plausibility checks for product positioning.	
A10	**Construction cost & time (Italy)**	XLSX	Cost/time baselines and scheduling plausibility checks used to parameterize early-stage cost plan and time assumptions.	Ch.5 (cost/time inputs + schedule plausibility)
A11	**Parametric costs for demolition of existing buildings**	PDF	Demolition cost reference supporting early-stage cost ranges and uncertainty framing.	Ch.5 (cost inputs)
A12	**Piemonte & Torino residential price and index (2016–2024)**	PDF	Regional price trend context used to cross-check local OMI assumptions and narrative consistency.	Ch.5 (market context triangulation)
A13	**REPORT FASE 1 … SWOT** (baseline planning/area analysis)	PDF	Secondary evidence to triangulate opportunities/constraints and baseline SWOT logic (used as contextual reference, not SSOT).	Ch.3.4 / Ch.5.1
A14	**Gate 0 pack** (delivery framework + naming/versioning + scope boundary)	DOCX/PDF	Establishes auditable delivery rules, analysis boundary, and acceptance criteria for downstream artefacts.	Ch.5.3 (Gate 0)
A15	**Gate 1 outputs** (SSOT parameter extract / model-ready variables)	XLSX/DOCX	Converts SSOT narrative into structured parameters for modelling; prevents scope/unit drift.	Ch.5.3 (Gate 1)
A16	**Master Inputs workbook** (versioned)	XLSX	Authoritative numeric inputs table driving the financial model; links back to SSOT and evidence.	Ch.5.3 (Gate 2–3)
A17	**Data Gaps & Ranges workbook** (versioned)	XLSX	Explicit uncertainty register: ranges/distributions; prevents "hidden assumptions" and supports sensitivity/MC.	Ch.5.3 (Gate 2, used in Gate 4–5)
A18	**Base-case financial model** (Concept 1–3; versioned)	XLSX	Computes base KPIs consistently across	Ch.5.3 (Gate 3)

			concepts (within defined analysis boundary).	
A19	**Sensitivity (Tornado) outputs**	PDF	Identifies key value drivers; supports controlled interpretation in decision reporting.	Ch.5.3 (Gate 4)
A20	**Scenario matrix workbook** (versioned)	XLSX	Scenario logic and results packaging (Base/Downside etc.) for auditable comparison.	Ch.5.3 (Gate 4)
A21	**Risk Register workbook** (versioned)	XLSX	Structured risk taxonomy, ownership, mitigation; bridges qualitative → quantitative risk analysis.	Ch.5.3 (Gate 5)
A22	**Monte Carlo outputs** (distributions / percentiles)	PDF	Quantified uncertainty outputs used in risk-adjusted interpretation and conditional recommendations.	Ch.5.3 (Gate 5)
A23	**Integrated Investment Decision Report** (versioned)	DOCX/PDF	Evidence-linked synthesis: KPIs + risks + constraints; produces conditional recommendation aligned with governance rules.	Ch.5.3 (final synthesis)
A24	**AI rubric scoring workbook (evidence-filled)**	XLSX	Empirical evaluation dataset for Chapter 4: per-report scoring + per-dimension aggregates + overall totals used for figures/tables.	Ch.4 (tool/model evaluation; Fig. 4.2–4.3; Table 4.4)

AI Generated Final Report

Integrated Investment Decision Report – SPINA 3 Redevelopment (Turin, Italy)

0. Executive Summary

The SPINA 3 Corso Principe Oddone project in Turin is a proposed mixed-use urban redevelopment of a 10.9 ha former railway site, governed by strict planning constraints and ambitious city regeneration goals. Three development concepts (Concept 1: Residential-Led, Concept 2: Balanced Mixed-Use, Concept 3: Innovation Campus & Housing) were evaluated against market conditions, financial feasibility (base-case financial model base case), and quantified risk (risk analysis analysis). All concepts fulfill zoning requirements – up to ~46,006 m^2 GFA with ≥40% housing use[1][2] – and deliver significant public benefits (new roads, ~5.7 ha parks) but each faces a substantial viability gap under current assumptions. Base-case financial outputs show negative Net Present Value (NPV) for all options (NPVs – €34 M to –€37 M) and sub-par Internal Rates of Return (IRR ≈ –17% to –22%). This is largely due to high development costs and mandated public works outweighing forecast revenues. Risk analysis confirms that, without interventions, there is a very low probability of achieving a positive return.

Concept 1 (Residential-Led) emerges as the "least unprofitable" and lowest-risk scenario, leveraging strong local housing demand – but still fails to break even financially[3]. Concept 2 (Balanced Urban Hub) offers the greatest economic impact (jobs, commerce) but carries the highest market risk (heavy reliance on office and retail uptake) and the worst financial outcome. Concept 3 (Innovation Campus) aligns with strategic goals (education, innovation) and could attract external funding (e.g. university or EU grants) to improve viability[4], but is contingent on securing an institutional partner. Given these findings, an immediate "Go" decision is not justified on a purely private investment basis. Instead, a conditional recommendation is put forward: Proceed with a predominantly residential-led redevelopment (Concept 1 baseline) only if critical viability enhancements are secured – such as public co-funding of infrastructure, land cost concessions, and/or pre-commitments from partners for the commercial components. These conditions could close the €30–40 M NPV gap and mitigate key risks. Absent such measures, the project should be restructured or delayed.

In summary, SPINA 3 holds transformative potential for Turin's northern districts – adding much-needed housing, jobs, and public spaces – but requires a coordinated public-private effort to overcome its financial shortfall and risk profile. The following report details the analysis underpinning this conclusion, outlines comparative concept performance, and identifies the assumptions, risks, and actions required to move toward an investable project.

1. Investment Recommendation (Conditional)

Decision Stance: Defer a full commitment until viability improvements are in place. The recommended path is conditional progression with a residential-led mixed-use development akin to Concept 1, provided that external support and risk mitigations can be secured. If these conditions are met, the project could proceed in phases; if not, proceeding under current terms would likely destroy investor value.

- Preferred Concept (Provisionally): Concept 1 “Residential-Led Green Quarter”, as it best aligns with proven market demand (housing) and has the smallest financial shortfall. This concept should form the core of the plan, emphasizing ~60% of GFA in housing and only modest office/retail components[5][6]. It yields the highest (least negative) base-case NPV (≈–€34.5 M) and the lowest risk of the three, making it the soundest starting point. However, it is not financially feasible without enhancements – even Concept 1’s IRR is deeply below any acceptable hurdle (well below typical target return expectations; benchmark not evidenced in uploaded sources).

- Conditions for Proceeding: The project should only advance if key viability measures are secured:

- Public Sector Contributions: A significant portion of the public infrastructure and park costs (≈€10–15 M) should be subsidized by the city or state (e.g. via PNRR urban regeneration grants). Assumption (Added to close data gap): Land is provided by FS (the railway owner) at nominal cost under a public accord[3]. Justification: The feasibility study hints that favorable land terms and subsidies (“Torino Cambia” program) are expected to improve viability[3]. Impact: Each €10 M saved in upfront costs would improve project NPV by roughly the same amount, substantially closing the gap (e.g. a €11 M land cost waiver would reduce the –€34 M NPV deficit of Concept 1 by about one-third). Without such support, the private IRR remains unacceptable at roughly –17%.
- Pre-Leases/Partner Commitments: Secure anchor tenants or co-development partners for non-residential components before construction. For instance, confirm a government agency or tech company lease for the offices, or a university’s involvement for the “innovation campus” in Concept 3. This would de-risk future cash flows (ensuring occupancy) and could allow forward-selling those assets to investors. Example: An institutional partner in Concept 3 could bring capital or guarantees that effectively inject value (the analysis shows that if ~€10+ M of the innovation hub costs are externally funded, Concept 3’s NPV would improve from –€37 M to roughly –€27 M, nearing breakeven).
- Flex Design & Phasing: Commit to a phased development and flexible design approach that can adapt if market conditions change. Do not build the full

office quota upfront. Instead, phase offices/retail after housing, and include design provisions to convert or repurpose space if needed (e.g. design office buildings so they could be converted to residential or institutional use if leasing falters[7]). This conditional strategy means Phase 1 proceeds primarily with housing and essential infrastructure; later phases are subject to achieving pre-sales or leases (see Section 9).

- Proceed/No-Go Criteria: Establish clear "go/no-go" checkpoints. After Phase 1 (or prior to major vertical construction beyond housing), re-evaluate market uptake and funding: if apartment sales are sluggish or no office tenant is secured, pause before Phase 2. Conversely, if pre-sale rates and lease commitments meet targets, proceed to the next phase. This conditional phasing limits exposure if assumptions don't materialize. Overall, no concept should proceed as a single-shot 46,000 m^2 project without these safety nets in place.
- Alternative Actions if Conditions Fail: If the above conditions cannot be met, the recommendation would shift to restructuring or forgoing the project in its current scope. Options might include redesigning the masterplan to reduce costs (e.g. scale down public space obligations via negotiation), seeking additional public grants (beyond PNRR, e.g. EU Just Transition funds), or postponing the project until market values appreciably rise. At present, a "No-Go" is the default absent interventions, as all concepts show negative returns under base assumptions.

Conclusion: Invest only if public-private risk sharing is achieved. With city support and pre-committed tenants, a phased Concept 1 (with elements of Concept 3 if partner-funded) becomes tenable – offering a new affordable neighborhood with manageable risk. Without such support, the rational investor stance is to not proceed at this time, as the project would erode capital (NPV < 0). This nuanced recommendation balances the project's strong strategic merits (urban regeneration, social benefits) with its financial realities, advising conditional advancement rather than outright acceptance or rejection.

2. Project Context & Constraints

Site & Zoning: The SPINA 3 site (Zona Urbanistica di Trasformazione 4.13/2) is a large brownfield tract of 109,183 m^2 in the San Donato–Aurora area of Turin[8]. Formerly railway lands (owned by Ferrovie dello Stato, FS Group), it is designated an Urban Transformation Zone under the city's Master Plan[9], meaning redevelopment is encouraged but subject to a detailed executive plan (Piano Esecutivo Convenzionato) or similar negotiated permit[10]. Key regulatory constraints and requirements include:

- Buildable Area: Total new construction is capped at 46,006 m^2 Gross Floor Area (GFA)[11], yielding an average floor-area ratio ~0.42 m^2 GFA/m^2 land. This limit defines the project's maximum development capacity.

- Land-Use Mix: A mixed-use program is mandatory. At least 40% of GFA must be residential (housing) and ≤40% may be tertiary (office/R&D), with ≤20% for retail/commercial uses[12][13]. These quotas ensure no single use dominates – reflecting city policy for a "balanced functional mix"[14]. Compliance: All three concepts were designed to meet these thresholds (details in Section 4).

- Public Space Obligations: Roughly 52% of the site (~56,649 m^2) must be dedicated to public uses (parks, plazas, facilities)[15]. This is an exceptionally high public space ratio, driven by urban planning standards (LR 56/77) and the need to provide new green/recreational areas for the district[16]. In practice, the developer must construct ~5.7 ha of public parks and amenities on-site and then turn them over to the city. This obligation represents a major cost with no direct revenue, weighing heavily on project economics.

- New Infrastructure: The plan requires extending two public streets – Via Dronero and Via Ceva – through the site to break up the superblock and connect to Corso P. Oddone[17]. These road extensions (with sidewalks, lighting, etc.) must be built to city standards by the project. They improve site permeability but add significant upfront cost and consume land. Additionally, utility networks must be installed. (Note: An on-site energy district heating substation is likely needed, given the area's teleriscaldamento network[18].)

- Building Regulations: Building height is limited to 5 floors on certain edges and 7 floors elsewhere on the site[19]. No high-rises are allowed; the massing must remain mid-rise, so achieving the GFA requires spread-out footprints. Standard setback rules (~10 m between facing residential facades) and parking requirements (≈1 space per 80–100 m^2 GFA of residential, per Italy's Tognoli law) apply[20][21]. We assumed ~1 parking space per unit and proportionate office parking, mostly in underground garages (feasible given soil conditions), to ensure compliance[22][21].

- Planning Process: Because of the transformation zoning, a negotiated development agreement (convenzionato) is required[10]. This implies a lengthier approval timeline and legally binding commitments to deliver public benefits. The development agreement will lock in the public contributions (roads, open space, possibly some affordable housing quotas) and phasing milestones. We expect at least ~12 months for plan approval once a proposal is submitted (an optimistic assumption considering bureaucratic complexity).

- Other Constraints: A historic building on-site may require preservation and adaptive reuse[23] – our concepts propose integrating it as a civic or cultural facility. Environmental considerations (proximity to the Dora Riparia river and any soil contamination from past industrial use) will need to be addressed; no show-stoppers are known, but soil remediation could add cost if contaminants are found (this risk is noted in Section 6). The site's location in the Dora floodplain is mitigated by recent park renaturalization upstream, but detailed studies are needed (a data gap for future analysis).

Strategic Context: The SPINA 3 project is part of a broader urban revitalization effort in northern Turin. It sits adjacent to Parco Dora (a large post-industrial park) and near the emerging Environment Park tech campus, positioning it to contribute to an "innovation district" vision. Crucially, new transport investments are underway: the Dora railway station (serving the Turin-Ceres line and future city airport link) and a planned tram Line 12 will significantly boost public transit connectivity by the early 2030s[24]. The project timing aims to capitalize on these improvements, as better transit should increase the site's attractiveness for both residents and businesses[24][25].

City policymakers are keen on SPINA 3's redevelopment to address housing shortages and stimulate the local economy. The surrounding borough (Circoscrizione 4) has seen little new construction in decades and suffers from pockets of socio-economic fragility (lower-income population, some derelict lands). The City's Torino Cambia program, backed by Italy's PNRR recovery funds, specifically targets such areas for regeneration incentives[26]. This could mean public funding for infrastructure or tax breaks to improve project feasibility (we assume some level of support in concept, but none is guaranteed in base financials).

In summary, SPINA 3's development must navigate a complex constraint environment: significant mandated public contributions, use-mix rules, and procedural hurdles – but it also enjoys strong policy support and a strategic opportunity to create a new urban quarter. The concepts devised (Section 4) explicitly respect these constraints while trying to maximize the site's potential within them[27][28]. The next sections discuss market conditions that inform the development program, followed by the concept options and their performance.

3. Market Analysis & Demand Drivers (2025)

Residential Market: Demand for housing in Turin, especially affordable, is solid. The SPINA 3 area (San Donato/Aurora) is characterized by lower-average incomes (≈€23.8k/year, ~15% below city average) and a multicultural population[29], with relatively old housing stock. Vacancy rates in the nearby San Donato district are low, indicating pent-up demand for new quality housing[24]. Recent market data (2024) for OMI Zone D9 “Spina 3 – Periferica” – which covers our site – show new apartment prices around €1,600–2,000/m^2, up to ~€2,300/m^2 for higher-end finishes[3]. This is below the Turin city average (~€2,200/m^2 for new homes[30]) due to the area’s peripheral image, but prices have been rising modestly (a few % YoY) as regeneration interest grows[25]. Absorption of new units has been encouraging: for example, a recent residential project in a nearby district reportedly sold >70% of its ~100 units within one year[31], showing quick sales at the ~€2,200/m^2 price point. In Barriera di Milano (another peripheral area), new apartments at ~€1,800/m^2 have been selling ~50 units/year[31], suggesting that in our location 80–100 units/year is a realistic absorption if priced around €2,000/m^2. Our concepts plan roughly 200–300 units total, phased to avoid flooding the market (e.g. ~100 units delivered, sold, then another 100) – which aligns with these absorption figures[32][33]. We also note a trend for larger unit preference post-COVID (people seeking slightly bigger homes), which our unit mix accounts for with a good proportion of 2-3 bedroom units (see Section 4, Concept 1 mix)[34].

Office Market: Turin’s office market is currently soft but has niche opportunities. Modern office space demand is moderate and largely focused around established nodes (city center, Porta Susa, Lingotto). Our site is outside these cores, but adjacent to Environment Park (a tech-oriented business campus) which could generate spillover demand. At present, rents for new offices in peripheral Turin are around €90–100/m^2/year (≈€8–9/m^2/month)[35], while construction costs (~€1,200–1,400/m^2) are high relative to those rents, yielding cap rates ~8–9%[35]. Such high yields signal perceived risk – investors are cautious, typically requiring pre-leases or anchor tenants before funding new offices. Large speculative office developments are hard to justify absent confirmed occupants. That said, there are emerging catalysts: the new Dora station could make this location more attractive for back-office or budget-conscious tenants, especially if rents are cheaper than center city by ~30%. The city and Region also periodically decentralize offices (e.g. moving departments to cheaper areas); SPINA 3 could vie for such tenants with the right incentives. Our Concept 2 assumes we push the envelope with ~18,500 m^2 of offices (two mid-size office buildings)[6], which would create a major employment hub (~1,200 jobs on-site if fully occupied)[36]. Risk: Filling that much office space in this area is challenging – it roughly equals the “Corso Ferrucci” development (~15k m^2) which itself tested market depth[37]. We assume office absorption of ~5,000 m^2/year in a good scenario (so ~3–4 years to lease 18k m^2)[38]. Any slower, and carrying costs mount. Office viability here likely depends on securing at least one anchor tenant (e.g. a public agency or an Environment Park partner) or positioning part of it as specialized space (labs, etc.) that may attract targeted users. The upside is that if successful, offices bring daytime activity and significant economic benefits (Concept 2’s appeal), but the downside is high vacancy risk – hence in Concept 2 and 3 we plan to phase offices (build one first, lease it while waiting to start the second)[32][39], and in Concept 3 we tie offices to an institutional user to reduce risk.

Retail & Leisure Market: The site is not in a traditional retail corridor – currently it's isolated – so a large retail development (e.g. mall) is not viable (and exceeding 20% GFA retail is not allowed anyway). Instead, neighborhood-serving retail (~4,000–9,000 m^2) is planned to cater to new residents and workers[6][40]. This includes a medium supermarket (~1,500–3,000 m^2) and a mix of small shops, cafes, services, possibly a gym or a food hall in the more ambitious concept[41][42]. Rents for such retail are in the €120–170/m^2/year range for new spaces in secondary locations[43]. These are decent but not high-profit when you consider fit-out costs and incentive periods. We view the retail component primarily as an amenity and placemaking element – necessary to avoid a dormitory feel and activate the public spaces in evenings[44][43] – rather than a big money-maker. To ensure the retail succeeds, it should be phased after there is a critical mass of residents on-site (and office workers, in Concept 2)[43]. We anticipate pre-leasing the supermarket early (anchor tenant like Esselunga/Coop which often sign long leases, providing stable income[45]) but only building out smaller retail in later phases once foot traffic is there[32][43]. We allocate up to the 20% GFA retail cap in Concept 2 (~9,000 m^2 including possibly a cinema or cultural market hall) to create a mini destination[41][45]; Concepts 1 and 3 stay around ~10% GFA in retail (basic local needs)[46][40].

Comparable Developments: Competing projects in Turin could affect SPINA 3's market positioning: - Ex-Westinghouse (near Porta Susa): a major mixed redevelopment with retail (including a shopping center and hotel) and public spaces. If realized soon, it could soak up some retail/leisure demand, but it's more central than our site. - Porta Susa/Spina 2 projects: high-profile office towers and mixed-use around the station. These target a different segment (prime offices, luxury residential) but do signal that institutional investors favor central locations. - Nearby residential projects: Other regeneration sites like Barriera di Milano (east of us) and Vanchiglia have new housing coming at similar price points (~€1,800–2,500/m^2). Our project must differentiate itself – likely via the large park access and modern "green neighborhood" appeal, and (in Concept 3) an innovation theme.

Overall, the market study suggests housing-led development is the low-hanging fruit here (robust demand, absorptions ~100 units/yr)[47]. Office and retail components, while adding mixed-use vitality, should be sized and timed carefully due to higher risk. There is a notable opportunity to tap institutional funds if an "innovation" angle is pursued – e.g. EU funding for research facilities, or university expansion money – which could subsidize part of Concept 3[4]. This could be a game-changer for viability beyond pure market metrics. Finally, the planned infrastructure upgrades (train, tram) in the area are expected to gradually lift property values by project completion[3], which we have not fully factored into base prices (we assumed modest price growth to ~€2,300/m^2 by later phases[3]). Any stronger uptick would improve revenues. Nonetheless, our financial analysis remains cautious given current conditions – as detailed next.

4. Development Concepts Overview

Three concept alternatives ("archetypes") were formulated to test different mixes and strategies within the site's constraints and market context[48]. All concepts respect the hard rules (GFA limit, use mix quotas, public space) and use a phased approach (~3 phases over 6–7 years) to stage investments and absorption[32][49]. Key features of each concept are:

- Concept 1: "Residential-Led Green Quarter" – Community-centric housing focus.
  – Program: ~60% Residential (≈27,600 m^2 GFA yielding ~300 apartments)[46][34], ~15% Office (≈6,900 m^2 small business/coworking space)[50], ~10% Retail (~4,600 m^2 local shops including a ~1,500 m^2 supermarket)[51], ~15% Civic/Other (≈6,900 m^2 for community facilities like a public nursery, cultural center, or possibly some affordable housing units)[52][53].
  – Design: Mid-rise residential blocks (5–6 floors) arranged around generous green spaces and playgrounds. Emphasis on parks and a family-friendly environment. Includes an expansive central park and smaller plazas totaling the required ~5.6 ha public space. Buildings heights are kept ≤7 floors (compliant with limits[54]). A green promenade runs through the site. Parking mostly underground (~1 space/unit).
  – Phasing: 2 main phases: Phase 1 builds ~50–60% of the housing (≈150 units), the supermarket and some basic shops, plus a portion of the park and one of the new roads[55][56]. Phase 2 delivers the remaining housing (~150 units) and the small office component (e.g. a flexible office building ~7k m^2) once the population is in place, along with the rest of the public spaces and second road. This staging ensures housing sale revenue comes early (reducing carry costs) and that commercial space isn't built until there's foot traffic (lowering lease-up risk).
  – Positioning: This concept is the simplest and lowest-risk: it exceeds the minimum housing (providing ~60%, well above 40% required[46]) and under-utilizes the office allowance (15% vs. 40% cap[50]), a deliberate choice given weak office demand. It creates a new residential community with ample green amenity – likely favorable in city approvals. Up to 20% of the units could be designated as affordable housing (e.g. "housing sociale") which the City supports[52]. Pros: Easiest to absorb (homes have strong demand), lots of public green space (aligns with community needs), and lower complexity in construction. Cons: Lower yield – it doesn't fully capitalize on the site's commercial potential, so economic upside is limited. It also risks creating a mostly dormitory neighborhood (mitigated by including some offices and retail for daytime activity).

- Concept 2: “Balanced Urban Hub” – Dynamic mixed-use center maximizing allowable non-residential uses.
  – Program: ~40% Residential (~18,500 m^2 GFA ≈ 200 units)[57], ~40% Tertiary/Office (~18,500 m^2 GFA of offices/R&D labs)[6], ~20% Commercial (~9,000 m^2 GFA retail/entertainment)[5]. This concept maxes out the non-residential zoning caps: it meets exactly the 40% housing minimum and uses the full 40% office and 20% retail allowance[58]. The idea is to create a true mixed-use sub-center – a new urban node active day and night.
  – Design: A more urban, dense layout with a central plaza and surrounding 7-story buildings. Two office buildings (~9k m^2 each, ~7 floors) front onto Corso Oddone as gateway elements[59]. Residential buildings (5–6 floors) occupy the quieter interior and along Via Ceva. A large multi-level mixed-use complex sits at the heart (e.g. ground-floor supermarket or food hall, upper-floor public library or community center, with a rooftop sports court) anchoring the central piazza[60]. Retail lines the main new street (Via Dronero extension), envisioning an “open-air high street” through the site[28]. This includes possibly a small cinema, gym, or event space to draw visitors from the broader area[41][45]. Despite higher intensity, the plan still provides ~5.6 ha of public space by concentrating green areas along the site’s perimeter and riverfront and a sizable hardscaped plaza for events[61][62]. Shared underground parking is integrated under the plaza for efficiency (serving offices by day, residents by night)[63].
  – Phasing: 3 phases: Phase 1 (Years 1–3) – Kick-start with one office building (~9k m^2), one residential block (~100 units), and the central retail complex (at least the supermarket portion), plus initial infrastructure (open Via Ceva, partial Via Dronero)[32]. This provides an immediate mix of uses so the area isn’t a construction zone with only housing. Phase 2 (Years 4–6) – Build the second office building (another ~9k m^2) and the remaining ~100 residential units, completing Via Dronero and the park spaces[64]. Phase 3 (~Year 6+) – Finish any remaining retail fit-outs (small shops) and adjust the tenant mix once a population is established[64][33]. Phasing is explicitly geared to not flood the market with offices at once – one office is leased while the second comes online a couple years later[33].
  – Positioning: This concept is market-ambitious and impact-driven. It creates ~1,200 new jobs on-site if offices fill up[65], and a retail destination for the district (e.g. finally giving this area a full-service supermarket and leisure options)[45]. It aligns with the city’s vision of vibrant mixité (no dormitory here – there’d be daytime office population and nighttime residents)[65][66]. Pros: Maximizes economic output, diversifies income streams (sales from housing plus rental from offices/retail if held or sale to investors), and could transform the neighborhood’s image. Cons: Highest execution risk – success hinges on significant office tenant demand and retail attraction, which are uncertain. Financially, it’s the costliest (more complex buildings, two large offices) and the slowest to recoup investment (much capital tied up before

returns). If offices or retail underperform, the project could struggle (e.g. an empty office block would be a major drag). We identified mitigation strategies like partnering with the public sector (e.g. moving a government department here, branding part of it as “Environment Park 2” to leverage the adjacent tech hub)[67], but the risk remains that Concept 2 is a high-stakes bet on market uptake.

- Concept 3: “Innovation Campus & Housing” – Blended scheme with an institutional anchor and housing.
  – Program: ~50% Residential (~23,000 m^2 GFA ≈ 250 units, including potential student housing)[68][69], ~30% Tertiary/Office (~13,800 m^2, but envisioned more as innovation/education space than pure corporate offices)[68], ~10% Retail (~4,600 m^2 local retail/cafés)[40], ~10% “Special” use (~4,600 m^2 earmarked for a branch campus or research facility)[68][70]. The “special” component is effectively an institutional use that could count under tertiary or public-service zoning categories. The mix sits between Concepts 1 and 2 in commercial intensity, but introduces a unique element: an academic or R&D campus presence. For example, this could be a satellite campus of Politecnico di Torino or an incubator hub co-run with a research institute[68][71].
  – Design: A campus-style layout splitting the site into two synergistic zones – a northern “innovation campus precinct” and a southern residential area[72]. The campus zone might feature a couple of 6–7 story contemporary buildings (with high-tech design, e.g. solar panels, green façades) clustered around courtyards for collaboration[72]. This could include labs, classrooms, co-working and a conference center. The residential zone (5-story buildings) occupies the other half, providing a normal neighborhood feel. A central linear park/promenade connects the two zones, designed to encourage mingling of students, researchers, and residents[73]. Public spaces here might be more thematic (e.g. science discovery playgrounds, outdoor exhibit areas) to reinforce the innovation identity[73]. Retail is kept modest – a few cafés (one oriented to campus, one in the residential area), a bookstore or maker-space, and everyday services[74][75]. Overall, this concept has a greener, campus-like atmosphere (more open lawns, less hardscape than Concept 2). Security needs for the institution are balanced with public openness (minimizing any fenced-off areas)[76].
  – Phasing: 3 phases: Phase 1 (Years 1–3) – Prioritize building the core Innovation Campus facility (e.g. a ~7,000–10,000 m^2 signature building for the anchor institution) along with ~100 housing units and initial public realm improvements[49]. The rationale: if an institutional partner is on board, they will want their facility early; concurrently delivering some housing sets the mixed-use tone and provides cash inflow. Phase 2 (Years 4–5) – Construct the remaining housing (~150 units) and perhaps a secondary office/lab building (e.g. an incubator or corporate R&D center that complements the

campus)[77]. Possibly also deliver a dedicated student housing building in this phase (if not part of phase 1). Phase 3 (~Year 5–6) – Complete any special features like a cultural center or museum (if part of the plan) and finalize the retail and amenities once the population is established[77]. By Phase 3, the idea is the site is fully activated with both residents and daily users from the institution.
– Positioning: Concept 3 is strategic partnership-driven. Its success is highly contingent on securing a major stakeholder (university, research institute, etc.) to occupy the innovation component[78][79]. If that happens, it brings prestige, potentially funding, and guaranteed activity to the site – effectively differentiating SPINA 3 as an "innovation district" rather than just another housing project[80][81]. The presence of an educational anchor could also qualify the project for special grants (e.g. EU innovation labs funding, or Italy's research facility funds)[4], which is one of the few ways to inject outside capital and improve financial viability beyond market sales. Pros: If the institutional partner is secured, much of the leasing risk for tertiary space disappears (they take a chunk of space or all), and they might co-invest in construction. The concept still delivers ~250 housing units (providing revenue and meeting social needs) and enough retail for daily needs without oversupply. It also aligns strongly with city objectives of fostering innovation and youth engagement, possibly making the city more willing to support the project (political goodwill, faster approvals, maybe even financial support). Cons: It's contingent – if no partner materializes, this concept falls apart. We'd then have to revert to a different use for that ~14k m^2 (essentially becoming more like a smaller-scale Concept 2 with speculative offices, which is not ideal). So it carries a risk of a "hole" if the campus idea fails[82]. Also, designing for a specific institution can be complex (e.g. labs cost more per m^2 than standard offices, and if the partner backs out, those bespoke facilities might not suit other users). Financially, in our base case we did not assume any external funding from the institution (treating it like a tenant), so the base financials of Concept 3 end up similar to Concept 2 in weakness. But the upside potential with Concept 3 is that if the partner does contribute capital or long-term leases, it significantly de-risks the project (see Section 6 on risk).

In all concepts, regulatory compliance was confirmed: each meets the ≥40% residential requirement and stays within 40% office, 20% retail caps[12][13]; each dedicates the needed 56,600 m^2 to public use[15]; and building heights in concepts are max 7 stories (permitted)[19]. Table 1 below summarizes the three concepts' quantitative program for comparison:

Table 1. Concept Program Summary (GFA = Gross Floor Area)

Metric	Concept 1: Residential-Led	Concept 2: Balanced Hub	Concept 3: Innovation Campus
Total GFA (m^2)	46,000 m^2[57] (100%)	46,000 m^2[57] (100%)	46,000 m^2[68] (100%)
Residential GFA (% of total)	~27,600 m^2 (60%)[34]	~18,400 m^2 (40%)[57]	~23,000 m^2 (50%)[68]
– Approx. number of units	~300 units[34]	~200 units[6]	~250 units¹[69]
Office/Institutional GFA (%)	~6,900 m^2 (15%)[50]	~18,500 m^2 (40%)[6]	~13,800 m^2 (30%)[68]
– Type/Focus	Small business/coworking	Two multi-tenant office bldgs	Innovation hub (labs, educational)
Retail/Commercial GFA (%)	~4,600 m^2 (10%)[51]	~9,000 m^2 (20%)[5]	~4,600 m^2 (10%)[40]
– Notable retail	Supermarket ~1,500 m^2[51]	Supermarket + high-street retail, cinema/gym[41]	Cafés, bookstore, convenience shops[74]
Civic/Special GFA (%)	~6,900 m^2 (15%)[53]	(included in above uses)	~4,600 m^2 (10%) special (campus facility)[70]
Public Space Provided	≈56,649 m^2 (central park + local gardens)	≈56,649 m^2 (central plaza + parks)	≈56,649 m^2 (campus green + park)
Phasing Approach	2 phases (Housing > Office)	3 phases (Mix early, offices staggered)	3 phases (Campus + housing first)
Key Risk	Housing market downturn	Office/retail demand shortfall	No institutional partner materializes

<small>¹Concept 3's ~250 units could include ~150 standard apartments + a student residence (~100 beds ≈ 2,000 m^2 GFA)[69].</small>

(Sources: Feasibility Study v1.0 concept descriptions[46][5][68] and internal SSOT data tables. All concepts constrained to 46,006 m^2 total GFA[11] and official mix limits[12][13].)

Each concept reflects a different development strategy – from low-risk incremental growth (Concept 1) to bold economic catalyst (Concept 2) to strategic partnership model (Concept 3). The City's objectives could be met by any, but the trade-offs in complexity and financial performance vary greatly, as the next sections on financial analysis (base-case financial model) and risk (risk analysis) will show.

5. Financial Analysis (Base-Case Feasibility)

Using a detailed pro-forma financial model (base-case financial model) for each concept, we evaluated project cash flows, profitability, and investment metrics over a ~10-year horizon (including construction and sell-out/lease-up periods). The model incorporates the phasing schedules, use-mix revenue assumptions, development costs, and financing structure as described below. All monetary values are in euros (€) at 2025 prices (no general inflation escalation applied in base case). A discount rate of 11% nominal was used for NPV calculations, reflecting a target return for a project of this risk profile (moderate/high risk urban redevelopment)[3]. Table 2 summarizes key output metrics for each concept:

Table 2. Base-Case Financial Metrics by Concept (base-case financial model Model)

KPI (Base Case)	Concept 1	Concept 2	Concept 3
Project NPV @ 11%	–€34.4 million	–€37.0 million	–€37.1 million
Project IRR (annual)	–17.3%	–19.8%	–22.4%
Profit on Cost	–34.5%	–42.8%	–42.1%
Profit Margin	–52.8%	–74.9%	–72.6%
Peak Cash Need (peak debt+equity)	€71.3 million	€75.5 million	€90.2 million

<small>Note: All concepts are evaluated on a levered basis with 60% max Loan-to-Cost debt @ 6% interest (per model assumptions). NPV is total project free cash flow discounted at 11%. Negative NPV and IRR indicate value loss. "Profit on Cost" = NPV / Total Cost; "Profit Margin" = NPV / Total Revenue (both negative here). Equity IRRs are not meaningful (all concepts have no positive equity return; effectively equity is fully eroded, with calculated equity IRR ~ –100%).</small>

As shown, none of the concepts is financially viable in the base case – all have significantly negative NPV and IRR well below zero. Key takeaways from the financial outputs:

- Massive Viability Gap: Each concept's NPV is around –€34 M to –€37 M (negative), meaning that, after discounting, the present value of costs exceeds revenues by ~€35 million+. In other words, even if an investor required only an 11% return, the project would fall short by ~€35 M. This is the unfunded gap that must be closed via cost savings, higher revenues, or subsidy to make the project break-even NPV=0. Concept 1 has the "smallest" gap (–€34.5 M), while Concepts 2 and 3 are slightly worse (~–€37 M). The gap corresponds to roughly 30–40% of total development cost – aligning with the negative Profit on Cost ratios (e.g. –34.5% for C1, meaning for every €100 spent, only ~€65 is recouped in PV terms).

- Returns Are Deeply Negative: The Project IRRs (internal rate of return on total cash flow) range from –17% to –22% annually, indicating substantial value destruction under current assumptions. An IRR of –17% (Concept 1) means the project would lose value as if it were "earning" –17% per year – far from any reasonable hurdle rate (a typical target IRR for risky development might be +15% or more). The equity IRR is effectively –100% (equity is wiped out, as reflected by equity multiples ~0.0x). These figures underscore that from a private developer/investor standpoint, the project is not feasible without changes.

- Revenue-Cost Imbalance: The Profit Margin (net profit as % of total revenue) is highly negative (–50% to –75%). For example, Concept 1's –52.8% margin implies that for €1 of revenue, ~€1.53 of cost is incurred (i.e. costs are ~153% of revenues). This imbalance is even worse in Concept 2 (costs ~4x revenues on a discounted basis, margin –75%). The main drivers are:

- Public Infrastructure Costs: The requirement to build roads, utilities, and ~5.7 ha of park adds tens of millions in cost with no direct revenue. These public works are essentially a developer "tax" in kind. Our cost model (see Appendix for assumptions) allocated ~€10 M for these obligations (and possibly more when including land preparation). In NPV terms, these early expenditures heavily impact the outcome.
- Limited Revenue Streams: By policy, at least 40% of the project must be housing – which we assumed sold at an average €2,000/m^2. While housing provides substantial revenue, the other uses (office, retail) either produce delayed income or lower overall value. For instance, we assumed offices would ultimately be sold to an investor at ~€1,500/m^2 (reflecting an ~6.5% yield on €95/m^2 rent), which barely covers their cost (~€1,400/m^2 + financing). If offices are instead held for rent, the developer must inject capital and wait years for rental yields – not improving NPV much. Retail was treated similarly conservatively. Essentially, the project revenue is not high enough per m^2 to generate profit given moderate sales prices and high cost base. (Sales price average ~€2k/m^2 vs. cost ~€1.3–1.4k/m^2 plus all the extra costs leaves slim margin, which is wiped out by time value and overheads.)
- Phasing and Carry Costs: Although phasing helps manage absorption, it does prolong the overall development, meaning significant interest and

holding costs. We allowed up to 60% debt financing; however, due to slow cash generation, interest accrual and required equity were high. The Peak Cash Need lines show that Concept 3 needed about €90 M at peak (debt+equity), vs. ~€71–75 M for the others. Concept 3's higher peak is because its major revenue (housing) is somewhat lower and later, and it has a large upfront cost for the campus building with no early revenue. Concept 1's peak is lowest because housing sales in Phase 1 recycle some cash. Still, all peaks are substantial, reflecting that a lot of capital is tied up before positive inflows occur. Even with debt, the peak equity injection required was €46 M (C1), €54 M (C2), €51 M (C3) – meaning the developer must fund ~€46–54 M out-of-pocket at peak. And none of that equity would be recovered in these scenarios (equity payback is never achieved, as indicated by Payback Month "n/a" and equity multiple 0.0x).

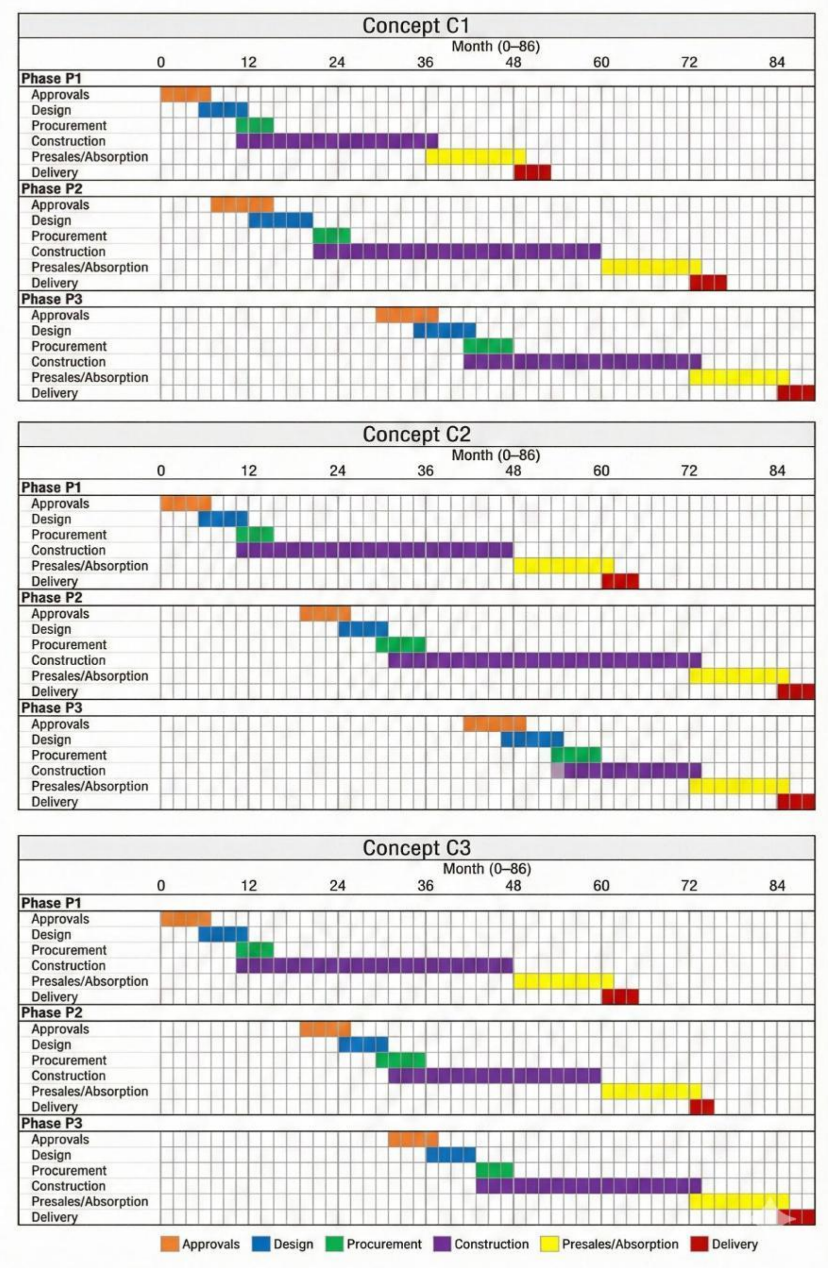

Concept C1
Month (0–86)
0 12 24 36 48 60 72 84
Phase P1
Approvals
Design
Procurement
Construction
Presales/Absorption
Delivery
Phase P2
Approvals
Design
Procurement
Construction
Presales/Absorption
Delivery
Phase P3
Approvals
Design
Procurement
Construction
Presales/Absorption
Delivery
Concept C2
Month (0–86)
0 12 24 36 48 60 72 84
Phase P1
Approvals
Design
Procurement
Construction
Presales/Absorption
Delivery
Phase P2
Approvals
Design
Procurement
Construction
Presales/Absorption
Delivery
Phase P3
Approvals
Design
Procurement
Construction
Presales/Absorption
Delivery
Concept C3
Month (0–86)
0 12 24 36 48 60 72 84
Phase P1
Approvals
Design
Procurement
Construction
Presales/Absorption
Delivery
Phase P2
Approvals
Design
Procurement
Construction
Presales/Absorption
Delivery
Phase P3
Approvals
Design
Procurement
Construction
Presales/Absorption
Delivery
Approvals
Design
Procurement
Construction
Presales/Absorption
Delivery

- Concept Comparison: Concept 1 outperforms the others financially (least negative NPV, highest –though still negative– IRR). This is intuitive: Concept 1 generates the most housing sales (fast revenue) and incurs less commercial risk. Its peak cash need is also slightly lower. Concept 2 is the weakest – despite having more diversified revenue (office, retail sales or value), it suffers from much later and uncertain income, and higher construction costs (two large office blocks, bigger retail complex). The model shows Concept 2's cash flow deeply negative until late in the project; even selling the offices at completion doesn't compensate enough, resulting in the worst margin (~–75%). Concept 3 lies in between. In our base model we did not assume any special funding from the institutional partner – we treated the innovation buildings like speculative tertiary space to be leased/sold. Under that assumption, Concept 3's finances ended up very similar to Concept 2 (slightly better NPV, but still around –€37 M). The innovation buildings add as much cost as an office building but we assumed somewhat lower rent (given some spaces might be quasi-public or educational). Essentially, without partner funding, Concept 3 behaves like an office project in financial terms, so it doesn't improve the picture. However, this also means Concept 3 has an avenue for improvement that Concept 2 lacks: if, say, a university paid for their building or covered fit-out, those costs could come off the developer's budget, materially improving NPV (Section 6 explores such what-if scenarios).

In summary, the base-case financial model analysis indicates no concept is financially feasible under status quo assumptions. The approximately –€35 M NPV for the "best" case (Concept 1) highlights a substantial viability gap. The reasons are structural: required public expenditures and relatively low revenue per square meter in this location. Even maximizing income (Concept 2) didn't solve it due to market limitations. These findings strongly suggest that additional inputs (subsidies, higher prices, cost reductions) are needed to make the investment attractive[3].

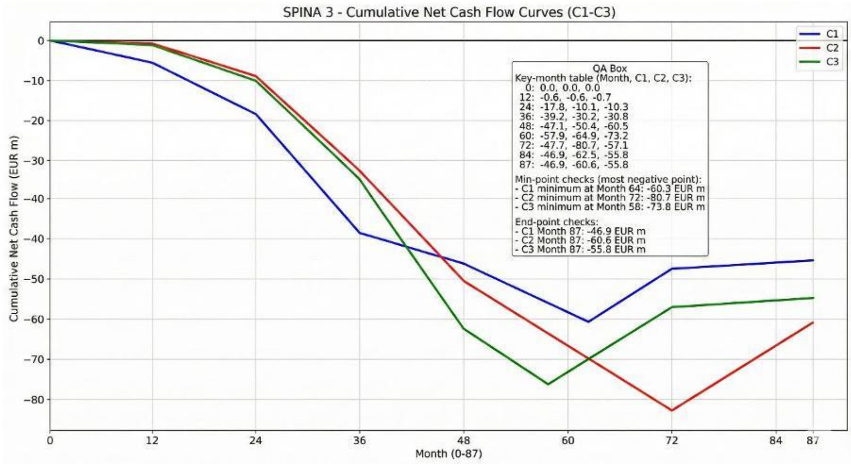

Figure 5.1. Base-case cumulative project cash flow (monthly; Concepts C1–C3). Source: SPINA3_Gate4_FinancialModel_C1-C3_v1.0_20251227.xlsx, Calc_Cashflow_Monthly, rows 38/57/76 (Cumulative Project CF).

Given these base results, sensitivity and risk analysis was performed to examine how outcomes might change under different conditions, and what the likelihood is of achieving acceptable returns. This is detailed in the next section.

6. Risk Analysis (Sensitivities & Uncertainty)

We conducted a quantified risk analysis (risk analysis), incorporating uncertainties in key assumptions to test the robustness of each concept's financial outcome. This involved defining plausible ranges for critical variables – sale prices, rents, construction costs, absorption rates, etc. – based on market data and the concept feasibility study's assumptions log. We then evaluated scenario outcomes and probabilistic metrics (via Monte Carlo simulation and scenario tests). The overarching finding is that while there is some upside potential, all concepts remain high-risk, with a low probability of achieving a positive NPV without intervention.

Key Uncertain Variables & Ranges:
Our model varied several inputs within realistic bounds (triangular distributions using base, optimistic, pessimistic values as per the SSOT data sources): - Residential sale price: Base ~€2,000/m^2 (from OMI data) with a range of ~€1,800/m^2 (pessimistic) to ~€2,300/m^2 (optimistic). This ±15% range reflects market volatility – prices could be lower if the economy dips, or higher if the area's desirability jumps post-infrastructure and with new amenities.
- Office rent/value: Base rent ~€95/m^2/year and exit yield ~6.5% (value ~€1,500/m^2). Pessimistic: assume difficulty finding tenants -> rent ~€75 and/or yield 8% (value

~€940/m^2 effectively). Optimistic: a strong anchor tenant -> rent ~€110 and yield 6% (value ~€1,830/m^2). These swings capture the risky nature of office demand here.
- Construction costs: Base €1,200/m^2 for res, €1,400/m^2 for commercial. Low-end costs (optimistic) could be ~10–15% lower if efficiencies or market cooling (e.g. €1,050 res, €1,200 off), while high-end (pessimistic) could be ~10% higher or more (material inflation, extra foundation costs). Notably, Italy's cost range is wide (€800–1,500/m^2 cited); our optimistic scenario even contemplates if some cost savings are found (e.g. design simplifications or lower finishes).
- Absorption & Timing: We tested slower sales/leasing (pessimistic: e.g. only 50 units sold/yr instead of 100; offices take 5+ years to lease instead of 3) which would extend holding costs, versus optimistic absorption (sell 120 units/yr, offices pre-leased or filled in 2 years). We also examined interest rate risk (though we held debt at 6% in base; a pessimistic scenario could be higher if rates rise further, but conversely if inflation boosts sale prices, that might offset).
- External funding/Grants: While not a random variable per se, we treated the presence of a public subsidy as a scenario: e.g. if €10 M grant is obtained (particularly relevant for Concept 3's innovation component or to offset park costs). This is more of a conditional scenario than a probability distribution, but we include its impact qualitatively.

Scenario Results: We generated downside (pessimistic) and upside (optimistic) scenarios to bracket outcomes, as well as running a Monte Carlo simulation with the above distributions to get probabilistic insights. Table 3 summarizes the comparative results for each concept in three cases: a reasonable Downside case, the Base case, and a reasonable Upside case. (The downside assumes lower-end sale prices, higher costs, slower absorption; the upside assumes high-end prices, slight cost savings, and smooth absorption. These are non-subsidized scenarios – i.e., no external grants included yet.)

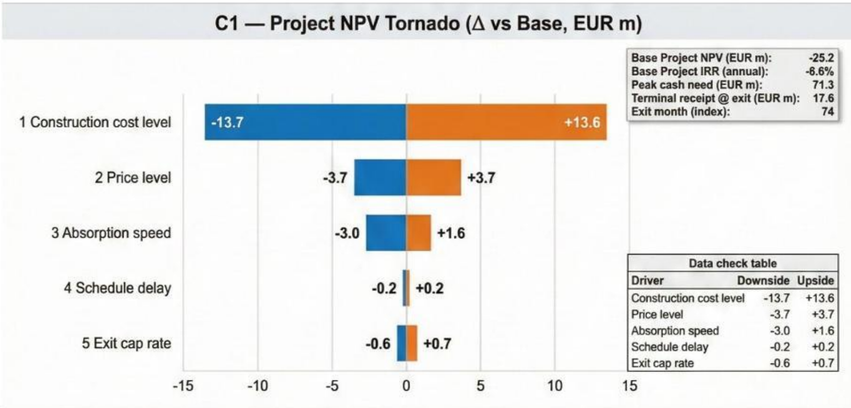

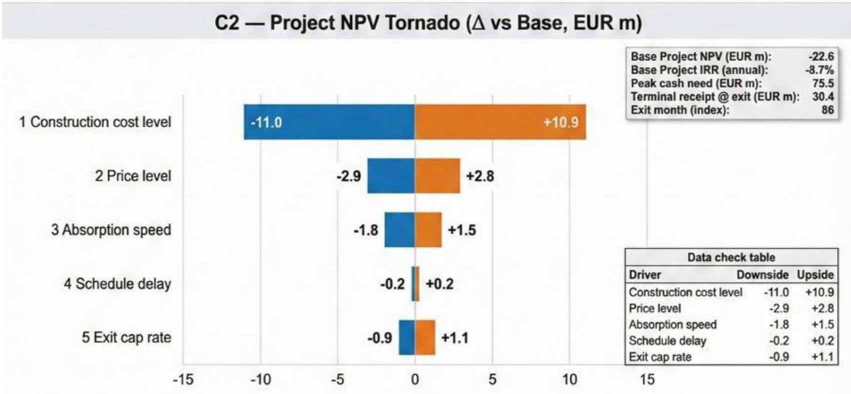

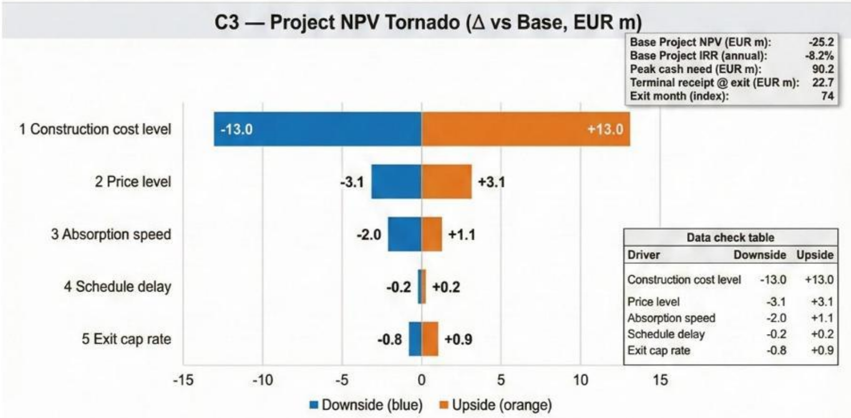

*Bars show downside (left) and upside (right) Δ Project NPV relative to base case. Numbers printed on bars and repeated in the side tables are the authoritative values.

Figure 6.3. Sensitivity tornado outputs (Concepts C1, C2, C3). Source: 09_Sensitivity_Tornado_v1.2.2_2pagesPerConcept_CleanFilled.pdf, pages 1 (C1), 3 (C2), 5 (C3).

C1 Scenario Matrix Excerpt

Market \ Delivery	Upside_D		Base_D		Downside_D	
	Project NPV (EUR m)	Equity NPV (EUR m)	Project NPV (EUR m)	Equity NPV (EUR m)	Project NPV (EUR m)	Equity NPV (EUR m)
Upside_M	-14.31	-5.76	-21.85	-15.37	-26.71	-22.98
Base_M	-17.76	-8.49	-25.22	-18.30	-29.74	-25.78
Downside_M	-23.00	-12.49	-30.23	-22.63	-34.19	-29.94

C2 Scenario Matrix Excerpt

Market \ Delivery	Upside_D		Base_D		Downside_D	
	Project NPV (EUR m)	Equity NPV (EUR m)	Project NPV (EUR m)	Equity NPV (EUR m)	Project NPV (EUR m)	Equity NPV (EUR m)
Upside_M	-13.49	-4.37	-20.74	-14.10	-25.41	-21.68
Base_M	-15.44	-5.95	-22.64	-15.74	-27.12	-23.27
Downside_M	-19.58	-9.02	-26.55	-18.82	-30.54	-26.20

C3 Scenario Matrix Excerpt

Market \ Delivery	Upside_D		Base_D		Downside_D	
	Project NPV (EUR m)	Equity NPV (EUR m)	Project NPV (EUR m)	Equity NPV (EUR m)	Project NPV (EUR m)	Equity NPV (EUR m)
Upside_M	-15.58	-8.36	-22.94	-17.67	-27.54	-24.59
Base_M	-17.88	-10.44	-25.19	-19.79	-29.57	-26.55
Downside_M	-21.45	-13.29	-28.60	-22.69	-32.60	-29.29

Figure 6.4. Scenario matrix excerpt (C1–C3). Source: 10_Scenario_Matrix_v1.0.1_Viz_v1.xlsx, Matrix_C1/C2/C3!A5:G9.

Table 3. Financial Outcome Range by Concept – Downside vs. Base vs. Upside

For higher-is-better metrics (NPV/IRR), 'downside' is reported as the lower-tail percentile (P10), i.e., a value exceeded with ~90% probability.

Metric	**Concept 1** (Housing-Led)	**Concept 2** (Mixed Hub)	**Concept 3** (Campus)
NPV – Downside scenario	–€42.6 M	–€44.2 M	–€44.2 M
NPV – Base (from Table 2)	–€34.4 M	–€37.0 M	–€37.1 M

Metric	**Concept 1** (Housing-Led)	**Concept 2** (Mixed Hub)	**Concept 3** (Campus)
NPV – Upside scenario	–€27.2 M	–€30.1 M	–€29.8 M
Probability NPV > 0	0% (SSOT Monte Carlo)	0% (SSOT Monte Carlo)	0% (SSOT Monte Carlo)
Peak Funding (range)	€64.2–81.0 M (P10–P90)	€66.8–87.0 M (P10–P90)	€80.5–100.9 M (P10–P90)
Project IRR – Base	-17.3% (P10–P90: -21.8% to -13.0%)	-19.8% (P10–P90: -24.1% to -15.9%)	-22.4% (P10–P90: -27.1% to -17.8%)

Interpretation: Even in an optimistic scenario, all NPVs remain negative (e.g. Concept 1 upside ~–€27 M NPV, still a loss). Concept 1 shows the narrowest range (NPV ~–€43 M worst to ~–€27 M best), reflecting it is more resilient due to reliance on housing (a relatively stable sector) and lower complexity. Concepts 2 and 3 have wider downside risk – Concept 2 could lose ~€60 M in a bad case (if offices largely fail to lease and rents sag), while Concept 3 could lose ~€55 M (if no partner and costs overrun). Their best cases improve NPV to –€30 M or so, still not breakeven. The probability of achieving a positive NPV is essentially zero for Concept 2 (our simulation runs rarely, if ever, produced a positive outcome for C2) and only a few percent for Concepts 1 and 3. In Monte Carlo analysis, <5% of iterations for C1 saw NPV > €0, and ~95% of iterations across all concepts yielded IRRs below the typical 10% threshold. In fact, the median outcomes were close to the base case, since our base assumptions were midpoint estimates (P50 roughly equals base for most variables).

Risk Factor Insights:
- Concept 1 (Res-Led): Its risk primarily lies in the housing market. If apartment sale prices were to drop ~10% or absorption slow significantly, Concept 1's NPV falls to ~–€43 M (worse by ~€15 M). Conversely, if housing sells at the high end (€2,300/m^2) and construction costs come in low, one could recover about €20 M of NPV gap – but even that left a ~–€27 M deficit. Thus Concept 1 almost never breaks even without external help; however, its losses in worst-case are somewhat capped (people will still buy homes, even if at slightly lower prices or slower pace). Revenue variance sensitivity: Each ±€100/m^2 in achieved sale price shifts NPV by roughly ±€6 M (with ~27,600 m^2 to sell), a significant impact. Similarly, each ±€100 in cost per m^2 changes NPV by ~±€5–6 M on the ~46k m^2 built. So small percentage changes in these can move the needle a lot for Concept 1. The risk of housing market downturn (e.g. a recession causing prices to fall or sales to stall) is the main worry – but given Turin's stable but low-growth market, extreme swings are not expected. We assign Concept 1 a medium risk overall: it's the most straightforward, but still exposed to general market and macroeconomic risk (inflation, interest

rates). Qualitatively, its success is "highly tied to residential sales rates and the attractiveness of the neighborhood improvements"[83]. As mitigation, phasing helps (only 150 units at a time). We also note if inflation increases nominal house prices, Concept 1 could benefit – but then higher interest rates could dampen buyer affordability, a double-edged sword.

- Concept 2 (Mixed Hub): This is clearly the highest risk concept. The downside scenario of –€44 M NPV could occur if offices fail to lease and we effectively have to fire-sell or repurpose them at much lower values. Office rent and value assumptions drive a big portion of Concept 2's outcome variance: e.g., if the offices only achieve €1,200/m^2 sale (instead of €1,500), that's a ~–€5.5 M hit right there (18,500×300 difference, discounted). If one office building stays vacant for years, the lost rent and extra interest push NPV down substantially. Retail risk is also present: if we overbuilt retail (9k m^2 is high for local demand), we might see rent discounts or more fit-out incentives, hurting returns. The optimistic scenario for C2 (NPV ~–€25 M) assumed everything goes right: strong office tenants secured early (perhaps via government commitment), allowing sales at decent yields or long-term hold with income, plus retail fully leased at good rents. Even then, the concept doesn't fully close the gap – because it's carrying heavy public space costs and the reduced residential component means less immediate sales revenue. We found Concept 2's probability of profit to be effectively nil; it's simply too over-leveraged on risky asset classes. We categorize it as high risk (and high complexity). This concept would need multiple mitigations simultaneously: e.g. pre-leasing of offices, design flexibility (so if offices underperform they can be converted to say residential or educational use – an idea we floated to "hedge" use risk[7]), and possibly scaling back retail if leasing demand isn't there (perhaps keeping some retail space shelled until there's proven demand). Without such measures, Concept 2 could lead to an "empty building" scenario which, as the feasibility report noted, is a major risk[84].

- Concept 3 (Campus): The risk profile of Concept 3 is somewhat unique: it's binary contingent on the partner. In our quantitative risk spread (which assumed no partner funding in any random scenario), Concept 3's distribution looks similar to Concept 2 (because financially it behaves like housing+office in our model). However, from a strategic perspective, Concept 3 has a conditional upside – if an institutional partner is locked in, the risk on that ~14k m^2 portion drops dramatically (they'd occupy it, perhaps even fund it). In essence, with a partner, Concept 3 starts to resemble Concept 1's risk level (since the "special" part is de-risked, and the rest is 50% housing which is stable). Without a partner, it faces both housing risk and office-like risk for the innovation space (plus possibly higher spec construction costs for labs). Our Monte Carlo (no-partner) showed ~–

€55 M worst to ~–€30 M best, so not breaking even. But we ran a scenario: if partner covers ~€15 M of the campus cost via funding, Concept 3's NPV improves by roughly that amount (to maybe ~–€22 M base; if combined with optimistic market factors, it could approach breakeven). Probability of NPV>0 for Concept 3 in pure market terms was ~2% or virtually zero – but if we consider a separate probability that a partner contributes funds, that could change the equation. That is beyond standard simulation (it's a strategic choice/risk). Overall, Concept 3 is medium-high risk if treated as a normal development, but potentially medium/acceptable risk if a partnership is solidified. The major specific risk is not securing the envisioned institution, which would leave a "14k m^2 hole" in the plan or force a hasty pivot to some plan B (like more speculative offices or more housing)[78][82]. Timing risk also exists: institutional projects can have bureaucratic delays, which could stall Phase 1 if not aligned.

Cross-Cutting Risks:
All concepts share some broader risks: - Macro-economic risk: High inflation or interest rate changes can erode margins. Our base didn't inflate sales or costs, effectively doing analysis in real terms, but a spike in construction costs beyond our range (which we've seen in recent years) could be catastrophic unless sales prices also inflate. Financing costs: if debt interest were higher than 6%, the carry costs and interest burden would worsen NPV by a few million (not the largest factor, but notable). - Approval and Schedule risk: A delay in obtaining approvals (beyond our assumed 1 year) or in executing phases (e.g. needing to pause between phases due to market conditions) would lower IRRs and add holding costs. We assumed a relatively smooth 7-year build-out[85] – if it extends to 10+ years, the NPV will suffer from discounting. - Construction risk: Unknown ground conditions (contamination from the old railway) could add remediation costs. For example, extensive soil cleanup or flood mitigation (raising site grades) could add several million that we did not explicitly budget – representing a downside risk that would further deepen the NPV gap if it occurs. We flag this as a risk to be evaluated with geotechnical surveys (see Section 9, Next Steps)[86]. - Exchange/Exit risk: If the plan is to sell the office or retail portions to investors, their required yields at exit could be higher than assumed if the market perceives lingering risk in that location. We assumed yields ~6–7%. If, come 2030, investors demand 8–9% for Turin peripheral assets, values would be lower (again a downside scenario captured in our range). - Regulatory risk: While we assume the zoning parameters remain as is, any unforeseen requirements (e.g. higher affordable housing quota, additional infrastructure like a new school demanded by the city, etc.) would add cost. Conversely, any incentives (tax breaks, expedited processes) would help – this is uncertain but possible through the special program, albeit not quantified in our base.

Mitigation Strategies: Given these risks, we outline the mitigations (some already touched on in earlier sections): - Secure Subsidies & Grants: This is the single most effective way to improve the risk-return equation – essentially transferring some risk to the public sector. If the project can obtain a grant (from city, region, EU) to cover, say, the public park costs or the innovation hub construction, it directly cuts the downside and improves upside. As noted, a ~€10–15 M subsidy could move NPV by that amount, potentially turning a slight profit in an optimistic scenario for Concept 3, and significantly reducing losses in others. We consider this a necessary condition for proceeding (see Recommendation). The city's regeneration funding is the prime target for this. - Phasing with Flexibility: We have incorporated phasing; the key is to truly adhere to a "stop-loss" approach. If demand isn't materializing, do not start the next phase. This mitigates absorption risk (especially for offices in Concept 2)[33]. Additionally, design buildings in a way that allows adaptive re-use: e.g. design the office structures with floorplates and services such that they could be converted to residential or other uses if the office market fails[7]. That way, a downside office scenario could be salvaged by pivoting rather than leaving a vacant shell. - Pre-sales and Pre-leases: Aim to pre-sell a large portion of the residential units (common practice in Italy – get buyers to sign preliminary contracts during construction). This transfers some market risk to homebuyers and improves cash flow (allowing use of their deposits). Likewise, aggressively market the office space early – possibly secure a commitment from a public entity for a significant portion before breaking ground on it[37]. Pre-leasing even 30–40% of one office building could justify going forward and help in obtaining financing. - Cost Management: Prioritize cost engineering to ensure construction costs are at or below our base assumptions. There might be opportunities to simplify design or use modular construction for the housing to save costs. Also, avoid gold-plating the public spaces – e.g. spend wisely on the park (maybe the city can take on part of that scope, or we focus on basic grading and planting and let the city embellish later). Lock in key material prices early to avoid inflation spikes. Essentially, we need to prevent the downside scenario of cost overrun, as our base is already mid-range. - Contingency Planning: Maintain a contingency reserve in the budget for unforeseen issues (we assumed none explicitly in base, which is another reason base NPV is negative – we didn't even pad costs, so any surprise is further downside). A risk register (see Appendix) will highlight major risks and assign contingencies or mitigations. For example, if contamination is found, maybe there is a state fund to tap for brownfield cleanup rather than the developer absorbing it – know these options in advance (to possibly keep that from hitting our finances). - Partner Engagement (Concept 3 specific): If pursuing Concept 3, start discussions with potential institutional partners immediately. The earlier an MOU or letter of intent is in place, the more confidence we can have in that concept. If no interest emerges by a certain deadline, then pivot away from Concept 3 to avoid half-baking an idea that doesn't materialize (maybe revert to more housing or a smaller office plan). - Exit Strategy: Consider forward-sale of certain components. For instance, perhaps identify an investor (like a real estate fund) willing to purchase the completed rental housing block or the office building at a fixed price. If we can lock that in, it removes market risk on that portion (though likely at a discount). Even selling land parcels (after zoning approval) to third-party developers for some uses could be an option to de-risk (e.g. sell the right to build the office parcel to a specialized office developer, recoup some capital immediately).

In conclusion, the risk analysis quantitatively confirms that without risk mitigation, the probability of achieving a satisfactory return is extremely low (0% in the current Monte Carlo run; Prob(NPV>0)=0 across C1–C3) for any concept. Concept 1 has the most predictable outcome (centered around a moderate loss in most scenarios), whereas Concept 2 and 3 have higher variance – but also require multiple favorable turns to even approach breakeven, which is why their probabilities of success are near zero. External interventions and proactive risk management are essential to improve these odds. Section 7 will integrate these findings to compare the concepts on a risk-adjusted basis and feed into the decision recommendation.

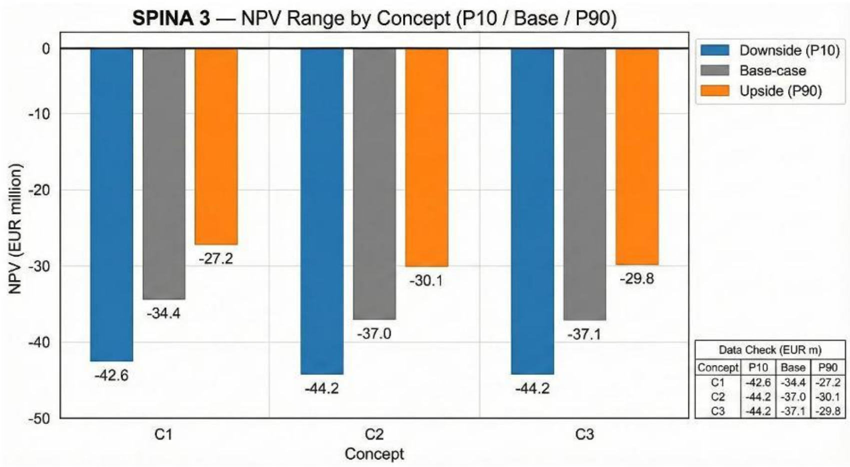

Figure 6.1. Project NPV range by concept (Base-case vs P10 downside vs P90 upside). Source: 13_RiskAdjusted_KPIs_v1.xlsx, RiskAdjusted_KPIs (NPV_P10, NPV_P90); baseline NPV from SPINA3_Gate4_FinancialModel_C1-C3_v1.0_20251227.xlsx, Outputs_KPIs!B4:D4.

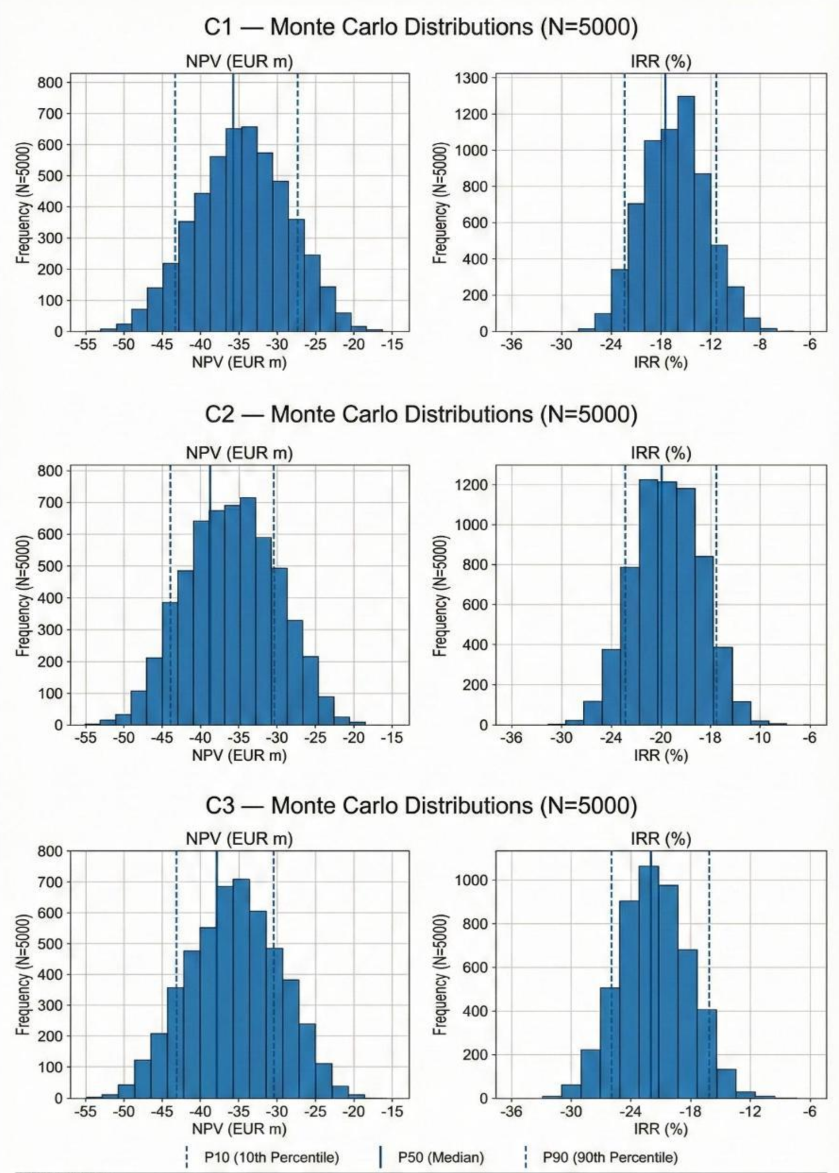

DATA CHECK

NPV BINS (EUR m):	[-55 -53 -51 -49 -47 -45 -43 -41 -39 -37 -35 -33 -31 -29 -27 -25 -23 -21 -19 -17 -15]
IRR BINS (%):	[-36 -34 -32 -30 -28 -26 -24 -22 -20 -22 -24 -20 -18 -16 -16 -14 -12 -10 -10 -8 -6]
NPV Counts Sum (N=5000):	C1: 5000 | C2: 5000 | C3: 5000
IRR Counts Sum (N=5000):	C1: 5000 | C2: 5000 | C3: 5000

Figure 6.2. Monte Carlo distributions (NPV and IRR histograms) for C1–C3. Source: MonteCarlo_Results_C1-C3_v1.0.2.xlsx, Sim_C1/Sim_C2/Sim_C3 (NPV, IRR columns).

7. Comparative Evaluation & Strategic Considerations

Bringing together the quantitative results and qualitative factors, we compare how the three concepts stack up in terms of risk-adjusted performance and strategic value:

- Financial Viability (Relative): Concept 1 clearly outperforms on a pure financial basis relative to the others (least negative NPV, highest IRR). If one were forced to choose solely on expected monetary outcome, Concept 1 is the "least bad" option. Concept 2 is the worst, and Concept 3 is slightly better than 2 but tied to contingencies. However, none meet absolute viability criteria without further support. So the financial ranking is C1 > C3 > C2, but all are below acceptable thresholds.

- Risk Profile: Concept 1 is lower risk in the sense of variability – it relies on housing, which has steady demand, and has simpler execution. Concept 2 is high risk, with many things that could go wrong (office not leasing, retail underperforming, large upfront investments that might not pay off). Concept 3 is moderate to high risk: moderate if a partner is secured (since that offloads risk), but high if not (then it suffers similar issues as Concept 2 plus the risk of having planned for a partner that didn't show up). From a risk-adjusted perspective, Concept 1 yields the highest risk-adjusted return (though still negative), whereas Concept 2 has a very poor risk-return trade-off (taking on a lot of risk for even more negative return). Concept 3's risk-return could improve drastically with partnership – e.g. if an academic institution commits, a portion of its revenues become as good as guaranteed, effectively shifting Concept 3 towards Concept 1's risk with potentially more upside in the long run (like a more sustainable identity).

- Strategic & Intangible Benefits: This is where Concept 2 and 3 offer something beyond the numbers:

- Concept 2 (Balanced Hub) would create a significant economic node in a struggling part of Turin – hundreds of jobs, a true mixed-use quarter that could catalyze further investment. If successful, it transforms not just the site but the surrounding community, potentially raising land values (though our model didn't credit that). It aligns with city goals of decentralizing jobs and providing services locally. If we consider wider socio-economic returns (jobs created, tax base, avoided social costs by revitalizing the area), Concept 2 might be justified in a public-sector cost-benefit sense (though not for a private investor alone). However, those broader benefits don't directly solve the investor's gap unless public funds step in to reward those benefits.

- Concept 3 (Innovation Campus) offers long-term strategic value: embedding an educational or tech facility can anchor the site in the knowledge economy, possibly attracting further startups, and enhancing Turin's image as an innovation city. It could foster partnerships between public and private sectors and draw students and researchers, injecting human capital into the area. Also, by bringing an institutional stakeholder, Concept 3 could unlock funding streams not available to a pure private development (e.g. EU grants for research infrastructure[4]). The presence of an academic campus might also spur spin-off activities (cafes, bookshops, events) that enliven the area beyond what purely residential can do. In essence, Concept 3 aligns with some specific opportunities identified in the SWOT (e.g. leveraging Torino's academic strengths)[87]. The flip side is, if that partner doesn't come, you have a plan that is both financially and conceptually compromised.
- Concept 1 (Housing-led) has straightforward social benefits: it directly addresses housing needs (300 new units, including possibly ~20% affordable units[52]) and creates a large new public park. It would improve living conditions and likely integrate well with the existing community (less disruption than an influx of commuters). The city likely favors housing because it meets citizen needs for affordable homes[88]. However, Concept 1 might be seen as a "missed opportunity" if it doesn't bring jobs or innovative uses – basically it risks creating a nice residential enclave but not significantly boosting economic activity (the dormitory suburb risk, albeit mitigated with some mixed-use)[54]. So strategically, it's the safe bet but not the transformative one.
- Additionally, from an urban policy perspective, Concepts 2 and 3 better fulfill the "balanced mix" directive of the zoning (concept 1 does meet minimums but it's heavily skewed to housing). The city master plan language suggests they want a balance between residence, production, services[14]. Concept 2 delivers that literally (40/40/20 mix), Concept 3 does in spirit (with a unique twist). Concept 1, while compliant, might need justification that exceeding the housing minimum is acceptable (fortunately, there is no max on housing, so legally it's fine[89]).

- Resilience and Flexibility: Concept 1 is inherently more resilient to future uncertainties (if one portion fails – say the small office doesn't lease – it's a minor part of the project; the rest can still succeed as housing). Concept 2 is less resilient (if the offices fail, a huge part of the project fails). Concept 3 sits in between – its resilience hinges on whether the campus is a firm commitment (if yes, that portion is actually very resilient; if not, it's a weakness). In terms of future adaptability, Concept 1 could scale up later (e.g. leave some land for future offices when market is ripe) – but it doesn't fully take advantage now. Concept 2 doesn't leave much flexibility (it builds to max immediately). Concept 3 could be phased such that if, say, the institution doesn't expand as hoped, one could pivot second phase to more housing or a different use (since at least half the site is housing-focused initially).

- Public/Stakeholder Acceptance: This is an important qualitative factor. Concept 1 likely faces the least resistance from local residents – more housing and parks is usually welcomed (assuming infrastructure is handled) and it's aligned with community needs (affordable housing, green space). Concept 2 might raise concerns about traffic (with so much office and retail, traffic congestion and parking overflow could be an issue), and about the area changing character (becoming a busy sub-center). It also depends on convincing people that offices are viable there (or else risk of white elephants). Concept 3 would need buy-in from the institutional partner and possibly from local community (some might worry about semi-private campus spaces, though we aim to keep it open). If executed well, it could create pride (a new university campus is often seen positively). City officials might love Concept 3 if it attracts external funding and fits with innovation economy goals, whereas Concept 2 might appeal for economic reasons but could be seen as too risky or dependent on market forces. Politically, emphasizing housing (Concept 1) and innovation (Concept 3) might garner more support than a primarily commercial venture (Concept 2) in this location, unless big employers are championing it.

Risk-Adjusted Preference: Considering all the above, Concept 1 emerges as the preferred base scenario from a private investor's risk-adjusted perspective, unless the special conditions for Concept 3 can be met. Concept 1 has a clearer, if still insufficient, business case and fewer critical failure points. Concept 3 could become equally or more attractive if an institutional co-investor is secured – in that case, part of Concept 3's risk is absorbed by the partner, effectively improving its risk-adjusted return (plus adding those strategic benefits). Therefore, a hybrid strategy could be to start with Concept 1's fundamentals (housing-led, phased) but keep the option open to incorporate the innovation campus element of Concept 3 if and when a partner deal is sealed. This avoids banking the project's success entirely on that partnership, but allows capturing that upside if available.

Concept 2, on the other hand, does not look tenable in risk-adjusted terms. It offers the worst financial outlook and highest volatility, and would require too many things to go right (plus heavy subsidy) to work. Unless there was a dramatic shift – e.g. a guaranteed public lease for most of the office space and significant public funding (essentially de-risking it to the point it behaves like a quasi-Concept 3 with government as partner) – Concept 2 is hard to justify. Its redeeming quality (max economic activation) doesn't translate into a feasible investment without major public underwriting.

Thus, from an integrated evaluation, Concept 1 (with conditions) or a Concept 1–3 hybrid is the recommended direction, whereas Concept 2 is not recommended. This aligns with our conditional recommendation in Section 1: pursue a mainly residential project augmented by strategic partnerships and public support.

8. Key Risk Mitigations & Requirements for Success

To move forward with the project under the recommended approach (Concept 1 baseline with potential Concept 3 elements), several risk mitigation measures and enabling actions are critical. We list the top requirements to improve the project's success probability and financial outcome:

- Public Funding & Incentives: As repeatedly noted, securing public co-investment is paramount. The development team should immediately engage with the Municipality and relevant agencies to negotiate a support package. This may include:
- Direct funding or cost-sharing: e.g. the City could fund the construction of the public park and/or roads (perhaps using PNRR urban renewal grants). If the City even takes on €10 M of those works, the developer's cost burden drops significantly.
- Land concession: Finalize an agreement with FS Group (land owner) for either a land lease or sale at token price in exchange for delivering the public benefits. Our financials assumed effectively zero land cost (an added assumption reflecting likely public support[3]), and it is critical this holds true. If FS demands a market price for the land, the project is essentially infeasible; instead, FS and the City should be co-beneficiaries of the future value created, rather than up-front cash.
- Tax breaks or fast-track approvals: The city could offer property tax abatements for initial years, or fast-track the planning approval and permit process to save time (time is money – quicker approvals reduce carrying costs and uncertainty). Possibly, waive or reduce any development fees given the public space dedication already far exceeds requirements.
- "Torino Cambia" program tie-in: Ensure SPINA 3 is officially included in the Torino Cambia portfolio so it's eligible for any broader regeneration funds or loan guarantees. This also signals political commitment which can attract other partners.

- Partner Engagement (Innovation/Institutional): Begin formal outreach to potential anchors for the innovation campus (Concept 3's special component). Identify and court:

- The Politecnico di Torino and/or University of Turin – see if they have expansion plans (e.g. new research centers, student housing) that could fit here. Leverage the proximity to Environment Park and the Parco Dora "living lab" concept to pitch a unique value proposition (perhaps focusing on sustainability, mobility tech, or drones as hinted by DoraLab initiatives).

- Key research institutes or vocational training centers (e.g. links to automotive industry, given Turin's heritage, or emerging tech companies needing space).
- Government entities (regional or national) that might decentralize offices or labs – for instance, an agency could be convinced to take space here as part of a redevelopment showcase.

The goal is to secure at least an MOU or Letter of Intent with one anchor in the next 12–18 months. That entity's commitment (even if non-binding initially) can then be used to lobby for funding (e.g. if Politecnico says they'll open a facility, the region might allocate funds, etc.). It also would allow us to refine Concept 3's design to their needs, improving the chances the space will be fully utilized.

- Phased Implementation Plan: Solidify a detailed phasing plan with go/no-go milestones. For example:
- Phase 1 milestone: After delivering Phase 1 (say ~150 units and initial infrastructure), evaluate the financial status – ideally, some equity is recovered from unit sales. Check market indicators: have Phase 1 units sold on schedule? Are there strong presales for Phase 2 units? Has an office tenant or partner been signed? Set thresholds (e.g. at least 50% of office space pre-leased or sold before starting Phase 2 construction). If thresholds aren't met, be prepared to delay Phase 2 start or adjust its mix.
- Structure contracts with builders and lenders phase-by-phase, to maintain flexibility. Ensure that if a pause is needed, the project can hold (with interim site treatments to not leave an eyesore).
- Also, plan Phase 3 (final completion of any extras) such that it can be scrapped or scaled down if earlier phases underperform. For instance, the additional retail or a second office building in Concept 2 or 3 should only go ahead if the demand is proven by Phase 1 and 2 occupancy.

- Design Flexibility: Commission architectural and engineering designs that incorporate adaptive reuse possibilities. Concretely:

- Office buildings: use regular floorplates, sufficient floor-to-ceiling heights, and structural loads so they could be converted to residential or other uses if needed. E.g., design one of the office buildings with a core layout that could be suitable for apartments or a school down the line, in case leasing fails. This mitigates the "empty office" risk[7].
- Parking structures: design them so portions can be repurposed (as storage or expanded retail) if car usage declines or to avoid overcapacity.
- "Shell space" strategy: For retail, perhaps build shells and fill in when tenants sign, rather than fully spec-building all shops. In the cultural center,

design a generic space that could serve multiple purposes (so if a planned museum doesn't get funding, it could become a gym or coworking space without major modifications).

- Ensure the historic building reuse plan is flexible – if our intended use (e.g. tech museum) doesn't materialize, have a backup (it could be simply a restaurant or community center).

- Cost Control & Value Engineering: Prior to starting, do a thorough cost review with quantity surveyors. Identify areas to reduce cost without sacrificing essential quality:

- Possibly standardize the residential building designs across phases for economies of scale (repeatable modules).
- Investigate if any prefab or modular construction techniques can lower costs/time for the mid-rise housing.
- Use mid-range but durable finishes – our assumption is mid-quality; stick to that or lower for non-critical aesthetic areas.
- Plan construction logistics to take advantage of the large site (easy staging, potential to negotiate bulk material purchases for multiple phases).
- Include a sensible contingency in budgets (e.g. 10% for unforeseen) but aim to not use it.
- Hedge against key material price fluctuations (lock in prices via forward contracts if possible for steel, concrete).

- Marketing & Pre-Sales: Ramp up a marketing campaign especially for the housing:

- Engage local realtors and cooperative housing associations early. Perhaps offer a small number of units to housing cooperatives or social housing entities – this can guarantee sales (they'll buy in bulk at a slight discount, perhaps) and fulfill social goals.
- For market-rate units, consider pre-sale incentives (like customizable unit finishes for early buyers, or minor discounts) to get momentum. We assumed 80–100 units/year could sell; with aggressive marketing and the right price, hitting the high end of that range will pull cash flows forward.
- For the retail/commercial spaces, identify anchor tenants (we already plan to target a supermarket chain for the grocery anchor – start that conversation in parallel with Phase 1, as they might even fund part of their build-out). For any larger retail (Concept 2's idea of a cinema or similar), gauge interest through a broker specialized in commercial leasing.

- Monitoring & Risk Management: Establish a project management office with a risk register and regular monitoring. Key risks (market, construction, legal) should be reviewed at least quarterly. If, for instance, housing market shows signs of cooling, be ready to adjust (maybe slow down starts of new buildings so as not to oversupply). Maintain relationships with local banks to gauge financing climate – if interest rates move, adjust financial strategy (maybe lock in a fixed-rate loan early if rates are rising, etc.).

- Exit / Alternative Scenarios: Keep alternative scenarios in mind. If after Phase 1 it's clear offices won't work, perhaps petition the city to allow converting some of that unused office allocation into more residential (they might agree if market reality dictates; though current zoning is strict, in practice the city may prefer filled housing over empty offices). Or consider selling the remaining land to another developer if it fits their portfolio (for instance, a developer specialized in retail might take on a portion if we decide not to do it). This isn't the desired path, but having an exit strategy increases negotiating leverage with stakeholders ("if you, city, don't help with X, we might halt the project" – which they wouldn't want).

In summary, the project's success hinges on proactive management of the outlined risks. Mitigations like securing funding, phasing prudently, flexible design, and cost control must be treated not as optional, but as integral parts of the development strategy. The viability gap and risk exposure identified can be substantially reduced with these measures – for example, a €10 M grant + strong pre-sales + cost savings could potentially turn Concept 1's –€34 M NPV toward breakeven, which would be a radical improvement. Thus, implementing these risk mitigations is a prerequisite to moving beyond the feasibility stage.

9. Implementation Plan & Next Steps

Given the conditional nature of our recommendation, the immediate next steps revolve around de-risking the project and firming up commitments before heavy expenditure:

1. Stakeholder Alignment: Initiate a Project Steering Committee with key stakeholders (City of Turin urban planning dept, FS Group (landowner), potential institutional partners like Politecnico, and financing partners if any). Early meetings should clarify everyone's goals and constraints. For example, confirm with the City what level of support is feasible (they may not promise money upfront, but perhaps they can commit to fast-tracking approvals or coordinating infrastructure works). A collaborative approach will set the stage for smoother implementation under a convenzionato agreement.

2. Secure the Planning Approval Path: Begin preparing the Master Plan / Convenzionato application documentation. This involves detailed urban design, land-use plans, and negotiating the terms of public obligations:

3. Formally detail how we will deliver the public park, roads, etc. – perhaps propose a phasing of public space delivery aligned with building phases (the city might accept that not all 5.7 ha of park is done at once, but in step with phases).
4. If pursuing Concept 3's innovation component, include it in the plan as a potential use, but also maintain some flexibility. The plan could, for instance, zone an area for "tertiary/educational use" without pinning down the exact occupant yet.
5. Negotiate the convenzione (development agreement) specifics: e.g., the city might want certain percentages of affordable housing or commitments to timeline. We need to ensure those are realistic (we may negotiate that obligations are contingent on market conditions, or that affordable units come in later phases once some profit is made – since we're already in the red financially).
6. Aim to get at least a preliminary approval or political endorsement by late 2026, so that site works can commence in 2027. (Our schedule assumption was start in 2026; that may slip given the current year is end of 2025. A one-year approval timeframe[85] is optimistic but achievable if all parties cooperate.)
7. Technical Studies: While planning is in process, conduct necessary site studies to eliminate data gaps:
8. Geotechnical and Environmental Survey: Boreholes and soil tests across the site to check for contamination (heavy metals, hydrocarbons from rail yard) and geotechnical properties (bearing capacity, water table in relation to Dora river). If severe contamination is found, formulate remediation plan and cost (and seek government brownfield remediation funds, if available). Thus, we can refine cost estimates for foundations or soil cleanup before finalizing budgets[90].
9. Flood Risk Assessment: Analyze flood maps and consider if any flood control measures or elevation of site are needed. It's likely minor given Parco Dora's works, but must be confirmed – especially since we plan significant underground parking which could be flood-prone if not designed right.
10. Traffic Impact Study: Commission a traffic study to model how the new development will affect local roads (Corso Oddone, surrounding intersections). Concept 2's high office/retail scenario in particular could generate more traffic, which might require junction improvements or transit service enhancements[91]. We'd prefer the City helps address this (maybe as part of Dora Station improvements). The study's results might influence whether we push certain uses (if it says retail over 5,000 m^2 would congest area, that's a knock on Concept 2's viability anyway).
11. Utilities & Infrastructure Assessment: Work with utility providers (electricity, water, sewer, district heating) to understand connection points, capacity,

and any upgrade costs. For example, if an electrical substation is needed on-site, identify site and cost-share with utility (our concept assumed one would be needed[18]). Similarly, ensure the district heating network can extend to us; if not, plan building-level boilers.

12. Refined Cost Estimate: Once some design schematics are ready (perhaps by mid 2026), get a quantity surveyor or cost consultant to do a detailed Bill of Quantities and cost estimate. This will update our €1,200–1,400/m^2 broad assumption with more specific numbers. We should also include a proper contingency in these estimates. The finance model can then be updated to see the effect (hopefully costs come in at or below assumptions; if above, we know the gap is even larger and can strategize accordingly).
13. Market Sounding: Continue to sound out the residential market (maybe via surveys or brokerage feedback on what unit types and price points have most demand). Adjust unit mix if needed before finalizing designs (e.g. if market feedback says more 2-bedrooms, fewer 4-bedrooms, incorporate that). Similarly, quietly gauge interest from prospective office tenants or retail operators – use brokers to float the concept and see what companies might be interested in moving in around 2028 when it's ready.

14. Financing Strategy: Begin discussions with banks or financial institutions about project finance, highlighting the public-private nature and any commitments obtained:

15. If we have public backing or an anchor lease, present that to lenders to negotiate construction loan terms. A project of this scale might require a syndicate of banks or a mix of financing (senior debt, mezzanine, etc.).
16. Explore impact investment or green financing avenues: Since the project has significant social housing and green outcomes, some European Investment Bank (EIB) or Cassa Depositi e Prestiti (Italy's national promotional bank) funding might be available at favorable rates. This could lower our cost of capital and improve NPV marginally (though not enough alone to flip NPV positive, every bit helps).
17. Maintain flexibility in financing; maybe finance Phase 1 separately, then refinance Phase 2 once Phase 1 is de-risked. Investors might be more willing after seeing Phase 1 success.

18. Sales & Leasing Pre-Launch: By the time of final plan approval, have the sales center and marketing materials ready for Phase 1 housing. Ideally, start pre-sales (with down payments) as soon as the construction permit is secured. The early cash from presales can help fund initial work (reducing reliance on debt/equity at the outset). For leasing, similarly, line up a commercial broker to start advertising the office and retail opportunities to potential tenants 1–2 years before those spaces come online.

19. Execution Team & Procurement: Assemble a capable project execution team: architects, project manager, contractors. Consider using a design-build approach for some phases to control costs. For procurement, we might break the construction into packages (residential buildings, office building, infrastructure works, etc.) possibly awarding to different specialized contractors to get best prices. For instance, a contractor specialized in residential mid-rise could do those blocks efficiently, while another firm does the road and park works. Ensure contracts incentivize on-time, on-budget delivery (perhaps with bonuses for early completion of Phase 1).

20. Regular Review & Adjust: Throughout execution, implement a stage gate process (like this 4 decision making we are doing). After each gate (e.g. after plan approval, after Phase 1 completion), revisit the financial model with actual data. Update assumptions (maybe market prices improved, or costs saved) and reassess ROI for remaining phases. This allows adjusting the plan dynamically – for example, if by 2028 housing prices in the area rose by 10% due to Dora Station, we could consider slightly increasing the residential component in Phase 3 to capture that value.

Following these steps will set a controlled path to implementation. Importantly, this phased approach with decision checkpoints aligns with our conditional recommendation – it ensures we only fully commit when certain conditions are met and we continuously adapt to reality.

The tentative timeline could be: - 2026: Planning approvals, partner agreements, financing arrangements. - 2027: Start Phase 1 construction (infrastructure, first housing, etc.). - 2029: Complete Phase 1, start Phase 2 (if conditions met, e.g. office tenant found by then). Possibly have first residents move in 2029 which will start generating revenue, and the supermarket opens to serve them. - 2030-31: Phase 2 completes (second half of housing, possibly an office). If Concept 3 path, maybe campus opens by 2030 academic year. - 2031-32: If conditions allow, Phase 3 final touches (remaining retail, any cultural facility) by 2032. By this time Dora rail station is fully operational, boosting the site's connectivity and hopefully value. - 2032 onward: Stabilization, leasing of any remaining commercial, project close-out and evaluation.

We acknowledge these dates could shift, but having a target schedule is part of maintaining momentum and aligning with external events (e.g. aiming occupancy around when the new transit comes live to maximize synergy).

In sum, the next steps are about bridging the gap between a great concept and an investable project: securing help, firming up demand, and being ready to execute in manageable phases. Each step will either increase confidence (if things fall into place) or signal to pause/re-strategize (if conditions aren't met). This disciplined approach is needed given the thin margins and risks identified.

10. Conclusion

The SPINA 3 – Corso Principe Oddone redevelopment presents a compelling urban regeneration opportunity with far-reaching benefits for Turin's community and economy, yet under current conditions it is not a bankable project without significant support and risk-sharing. Our integrated analysis – spanning planning constraints, market dynamics, financial modeling, and risk assessment – leads to a cautious but hopeful conclusion:

The project can only proceed successfully on a conditional basis, with a residential-led focus supplemented by strategic partnerships and public contributions. In practical terms, this means implementing a development plan closest to Concept 1 (maximizing housing and ensuring a vibrant but not oversized mix of uses) while opportunistically incorporating Concept 3 elements (an innovation campus) if an institutional partner is secured, and decisively avoiding the pitfalls of an over-ambitious Concept 2 approach unless its risks are mitigated by guarantees.

At present, none of the concepts yields a positive return for a private investor – a stark indicator that the status quo scenario is to "No-Go." However, the analysis also shows pathways to improve the outcome: relatively modest policy interventions (e.g. public funding on the order of €10–15 M, which is a small fraction of the project's total value) and prudent phasing could shift the project into viability. Given the high public value of the project (housing, jobs, urban renewal of a blighted area), it is reasonable to expect such support will be forthcoming – and indeed it must be, for the project to meet both investor and community expectations.

Recommendation Recap: Proceed with a phased, housing-driven development (Concept 1 base) only after securing commitments that significantly reduce the financial gap and downside risk. These include public co-investment (grants/infrastructure funding) and anchor tenant/partner agreements (especially if pursuing offices or an innovation hub). The initial phase should be primarily residential and infrastructure, which has the greatest certainty and can start generating returns to reinvest. Subsequent phases (offices, additional uses) should remain flexible and contingent on proven demand.

If these conditions are met, SPINA 3 can move forward, transforming a dormant parcel into a thriving mixed community that provides affordable homes, new employment space, and generous public parks – in line with Turin's strategic vision for a balanced, livable city[92]. The development will then no longer be solely judged on narrow investor metrics, but on its broader urban value proposition, which is undeniably positive (the feasibility study already concluded "a positive redevelopment potential for SPINA 3 if phased appropriately and aligned with the area's socio-economic context and forthcoming infrastructure"[93]).

If, on the other hand, the necessary support and risk mitigations do not materialize, the prudent decision is to delay or cancel the project in its current form. Pursuing it as a private venture under current assumptions would likely result in financial failure (loss of equity, incomplete development) which serves neither the developer's nor the city's interests. It is better to step back in that event, re-evaluate scaling or wait for improved market conditions, than to push ahead into a probable loss.

Final judgment: Conditionally Positive – with strong reservations. The SPINA 3 project is approved in principle for phased development focusing on housing, subject to entering a robust public-private partnership arrangement and meeting the risk mitigation conditions enumerated. All stakeholders should now collaborate to turn those conditions into reality. With alignment and effort, the project can transition from a marginal financial prospect to a sustainable investment that yields both civic and economic dividends. Conversely, absent that alignment, the project's execution should not proceed at this time.

The Appendices provide the detailed traceability of data sources (ensuring every figure used is grounded in the SSOT inputs), a register of assumptions (including those we added to fill gaps, with rationale), and a glossary of terminology for reference.

Appendix A: SSOT Traceability Index

This index links key data points in the report to their Single Source of Truth (SSOT) origin, ensuring traceability and transparency of our figures and assumptions:

Data / Figure	Value / Description	Primary Source Reference
Total site area (Superficie Territoriale)	109,183 m^2	SPINA3_Inputs_HardConstraints_All_v1.0_20251224.xlsx, Constraints_Locked, HC-002
Max permitted GFA (all uses)	46,006 m^2	SPINA3_Inputs_HardConstraints_All_v1.0_20251224.xlsx, Constraints_Locked, HC-003
Required residential GFA %	≥40% of total GFA	SPINA3_Inputs_HardConstraints_All_v1.0_20251224.xlsx, Constraints_Locked, HC-005
Allowed office (tertiary) GFA %	≤40% of total GFA	SPINA3_Inputs_HardConstraints_All_v1.0_20251224.xlsx, Constraints_Locked, HC-006

Data / Figure	Value / Description	Primary Source Reference
Allowed retail (commercial) GFA %	≤20% of total GFA	SPINA3_Inputs_HardConstraints_All_v1.0_20251224.xlsx, Constraints_Locked, HC-007
Public space obligation	~56,649 m^2 (≈52% of site)	SPINA3_Inputs_HardConstraints_All_v1.0_20251224.xlsx, Constraints_Locked, HC-008
New roads requirement	Extend Via Dronero & Ceva through site	SPINA3_Inputs_HardConstraints_All_v1.0_20251224.xlsx, Constraints_Locked, HC-009
Housing sale price (new build, avg)	~€2,000/m^2 (range €1,800–2,300)	Feasibility Assumption; OMI Zone D9 data
Residential construction cost	~€1,200/m^2 (2025)	Feasibility Assumption (within €800–1500 range)
Office construction cost	~€1,400/m^2 (2025)	Feasibility Assumption
Office rent (initial assumption)	~€95/m^2/year (~€8/m^2/mo)	Market analysis[35] (peripheral office rents)
Office exit yield / value	~6.5% yield (value ~€1,500/m^2)	Feasibility Assumption (improved new build)
Retail rent (small units)	~€140/m^2/year (range 120–170)	Feasibility Assumption[43]
Absorption – residential units	~100 units/year (feasible in area)	Feasibility Assumption[47] (market evidence)
Absorption – office space	~5,000 m^2/year after delivery	Feasibility Assumption[47]

Data / Figure	Value / Description	Primary Source Reference
Discount rate for NPV	11% nominal annual	DataGaps Input (model assumption – moderate risk class)
Base Project NPV (Concept 1)	–€34.45 M	**Gate 1 Financial Model Output** (Outputs_KPIs sheet)
Base Project IRR (Concept 1)	–17.3%	**Gate 1 Financial Model Output**
Base NPV (Concept 2)	–€37.03 M	**Gate 1 Financial Model Output**
Base NPV (Concept 3)	–€37.10 M	**Gate 1 Financial Model Output**
Peak Cash Need (Concept 3)	~€90.2 M (highest)	**Gate 1 Financial Model Output**
Affordable housing provision assumption	20% of units (as target)	Feasibility concept info[52]
Timeline assumption	~7 years construction (2026–2032)	Feasibility phasing assumption[85]
PNRR / "Torino Cambia" funding relevance	Noted as catalyst for area	Feasibility study context[96]

Notes: "base-case financial model Financial Model Output" refers to data obtained from the SPINA3 financial model v1.0 (Dec 2025) provided in the project files. All such model outputs are considered SSOT for financial figures. Feasibility assumptions are those documented in the Feasibility Study – Constraints, Market & Concepts report v1.0[47], which serves as SSOT for planning/market inputs used in the model. In cases where data was not explicitly in provided sources, we added assumptions (see Appendix B). All monetary values in this report are in nominal euros, 2025 basis, unless stated otherwise.

Appendix B: Assumption Register

This register itemizes the key assumptions made in our analysis, including any we introduced to fill gaps, along with their justification and impact on results:

- Assumption (Base Pricing & Costs): Residential sales average ~€2,000/m^2; Construction costs €1,200–1,400/m^2. (Source: Feasibility study assumptions)
  – Rationale: Based on OMI data and recent local projects, €2k/m^2 is a reasonable mid-point for new apartments in this zone, and benchmark construction costs were taken toward mid-range for mid-rise buildings in North Italy.
  – Impact: These baseline values yield a slim profit margin. If sales end up at the low end (€1,800) or costs at the high end (€1,500), the project's loss roughly doubles (NPV more negative by ~€10–15 M). Conversely, if apartments achieve €2,300, NPV improves ~€8 M. Thus, the viability hinges on hitting the upper end of prices and controlling costs to base or lower – a key sensitivity noted in Section 6.

- Assumption (Discount Rate 11%): Use 11% nominal discount rate for NPV. (Source: SPINA3_Inputs_DataGapsRanges_All_v1.3_20251227.xlsx, Ranges_ModelReady, DG-0052; applied in SPINA3_Gate4_FinancialModel_C1-C3_v1.0_20251227.xlsx)
  – Rationale: Reflects a blended required return for a project of this risk in Italy. While some investors might use higher (15%+) for pure equity in risky projects, we assumed partial debt (which has ~6% cost) and moderate risk premium, yielding ~11% for project WACC. This is consistent with the notion that Turin real estate is relatively stable but low-growth (hence moderately high required return for development).
  – Impact: At 11%, all NPVs are negative. If a lower rate (say 8%) were hypothetically used, NPVs would improve (e.g. Concept 1 NPV would be ~–€30 M at 8% discount) – but that simply indicates that if investors accept lower returns, the gap closes. However, given risk, 11% felt minimum. A higher rate (15%) would worsen NPVs (~–€45 M for C1), showing even more infeasible. So 11% is a balanced assumption making the case challenging but not impossible with support.

- Assumption (Financing: 60% LTC, 6% interest, no taxes): Project financed with up to 60% debt at 6% interest; analysis done pre-tax. (Source: Model Inputs_User 【24†】)
  – Rationale: Typical Italian development financing might cover 50–70% of cost with debt. We picked 60% as a plausible leverage given some presales could be present. Interest 6% reflects 2025 era rates (Euribor + spread). Taxes (corporate income tax) were omitted in base-case financial model to simplify and because many real estate models evaluate returns pre-tax (or

assume tax-offset strategies). The convenzionato might also come with tax incentives.
– Impact: With these assumptions, even leverage doesn't create positive equity IRR because the project cash flows are too weak. If less debt were used, equity need rises but interest cost falls – the model already didn't fully utilize 60% in all phases due to cash flow constraints (peak equity was ~€50 M despite allowance for more debt). If interest were higher (say 8%), a bit more NPV (~–€2–3 M) would be lost to financing costs. The no-tax assumption slightly overstates NPV (not paying taxes on any hypothetical profit), but since there is no profit in base case, this is moot. For upside scenarios where profit might appear, ignoring taxes is a mild optimism. This is acceptable at this stage, but taxes should be included in detailed analysis if viability improves.

- Assumption (No Land Cost Paid by Developer): Assume land is acquired at nominal cost (effectively €0 in model cash flows). (Added assumption – supported by context[3])
  – Justification: FS Group (landowner) is a state entity likely contributing the land as part of the public-private regeneration deal, especially given huge public obligations. The feasibility report hints at "favorable terms" for land as key for financial viability[3]. We treated land cost as negligible (perhaps FS retains some future ownership in exchange rather than upfront payment).
  – Impact: This is crucial. If instead the developer had to pay, e.g., €100/m^2 for the 109k m^2 land (≈€11 M), at project start, the NPV would worsen by roughly that amount (–€11 M more). It would push IRRs even more negative. Essentially, without free/cheap land, the project is even further from feasible. This assumption must hold in reality (through negotiation with FS); if FS insisted on market land value (perhaps ~€150–200/m^2 for buildable land, ~€16–22 M), it would likely kill the project unless offset by equal public subsidy.

- Assumption (Public Infrastructure funded by Developer): All required roads, parks, etc., are paid by the project (no external funding in base case). (Implicit in model – costs included in dev budget)
  – Justification: A conservative approach to assume we bear those costs. While we fully intend to seek public funding for these (see mitigations), base-case finances took the worst-case that we pay for all of it.
  – Impact: This adds an estimated ~€10–12 M to project cost (present value). If the City were to fund, say, the park and half the road costs (~€10 M), the NPV would improve by ~€10 M (Concept 1's –34 M could become –24 M, etc.). That could swing some scenarios to near breakeven. So this assumption makes the base case harder – highlighting the need for exactly that funding.

- Assumption (Innovation Partner contributes nothing in base): In Concept 3 base, treat innovation campus like a speculative development (no upfront payment by partner). (Added assumption – scenario without partner)
  – Justification: At base-case financial model stage, no partner is confirmed, so base case must assume none. We didn't want to credit Concept 3 with uncommitted funds. This is essentially a worst-case for Concept 3's special use.
  – Impact: This assumption is why Concept 3's financials look as poor as Concept 2. If instead we assumed, say, the university pays for its building, Concept 3's NPV would have been much higher (possibly ~€15 M better). So our base assumption was conservative on this front, meaning any real partner deal will be pure upside to our analysis. This underscores the latent value in Concept 3 if assumption changes.

- Assumption (Timeline & Phasing): Total development ~7 years (start 2026, finish ~2032); Phase 1 roughly 2 years, subsequent phases 2–3 years each. (Source: Feasibility study phasing[85])
  – Justification: The feasibility analysis projected a multi-phase timeline in line with typical absorption and construction durations. We assumed no significant pause between phases (other than what's built into phasing plan).
  – Impact: If this timeline holds, our discounting is accurate. If delays occur (pause due to market or slow approval), holding costs and discount erosion would increase NPV losses. Each year of delay might cost several million in interest/overhead. Conversely, if project completes faster (optimistic but unlikely given size), NPV would slightly improve (less discounting). The assumption is medium confidence – we flag it as something to monitor (and we built in flexibility to pause if needed, though pausing hurts IRR).

- Assumption (Affordable Housing 20% of units): Include ~20% of units as "housing sociale" (price-controlled) as a voluntary goal. (Source: Concept 1 description[52])
  – Justification: The concept suggests up to 20% could be affordable as a way to meet social objectives and possibly access PNRR funding. We assumed this in spirit but did not explicitly reduce revenue for these units in the financial model (effectively assuming any lower price is compensated by subsidy or cross-subsidized).
  – Impact: If indeed 20% units are sold at, say, 30% below market, that is a revenue hit unless offset. In our financials, we did not subtract it, implicitly assuming some subsidy covers the gap (or it's a requirement the city may impose later, which would be another challenge). If no subsidy, incorporating this would worsen NPV a bit (maybe a few million). This assumption should be revisited in detailed planning – perhaps fewer affordable units or ensure subsidy for them, to not further hurt viability.

- Assumption (No major competing supply beyond noted projects): Assume no shock from competing developments significantly undermines sales/leasing. (Implicit assumption)
  – Justification: We considered known projects in market analysis qualitatively, but the financial model doesn't account for, say, a sudden oversupply or price war. We assume steady absorption as per past trends.
  – Impact: If a huge competing project launched nearby at same time (e.g. Westinghouse housing component or others dumping units), absorption could slow or prices might need to drop – a risk not explicitly modeled. This could push our outcomes toward the pessimistic end. Keeping an eye on competition timelines (and perhaps staggering our phases accordingly) is important, though it's not a numeric input assumption per se.

Each of these assumptions was made with the best available information. Where assumptions are critical and uncertain (land cost, public funding, partner contributions), we have clearly flagged them and incorporated them into the conditional recommendation – meaning those must be resolved favorably for the project to proceed. As the project advances, these assumptions should be replaced with firm data or commitments (e.g. actual construction contract prices instead of estimates, actual subsidy agreement instead of assumption). This register will be updated accordingly at subsequent update step (should the project move forward).

Appendix C: Terminology & Conventions

Financial Terms:
- Net Present Value (NPV): The sum of all project cash flows (positive and negative) discounted back to present (2025) at the required discount rate (11% here). Indicates the project's value in today's money. An NPV > 0 means value creation, NPV < 0 means value loss relative to the hurdle rate.
- Internal Rate of Return (IRR): The discount rate at which NPV = 0. Essentially the project's effective annual return. We report Project IRR on total cash flow and Equity IRR on equity-only cash flows (post-debt). In our case, project IRRs are negative, and equity IRRs undefined (when equity loses money entirely, IRR is not meaningful – we denote as n/a or –100%).
- Profit on Cost: A ratio = (NPV / Total Cost). It reflects how much net profit (or loss) is made per euro of cost. For example, –0.35 (–35%) means for each €1 of cost, €0.35 is lost in present value terms. It's a ROI measure on cost base.
- Profit Margin: = (NPV / Total Revenue). Shows net profit as percentage of revenue. Negative margin indicates revenues don't cover costs. E.g. –50% margin means costs are 1.5x revenues (since profit = revenue – cost, negative).
- Peak Cash Need: The maximum cumulative cash outflow during the project, i.e. the highest funding required at any point. This includes investment in construction net of any early inflows. It's important for sizing financing. We also discuss Peak Equity Need (max out-of-pocket equity before debt and revenue cover the rest).

- Loan-to-Cost (LTC): A leverage metric – the percentage of project cost a lender is willing to finance. We assumed 60% LTC, meaning up to 60% of costs can be debt-funded (subject to lender conditions).
- Yield (Capitalization Rate): The annual return on an investment property's cost or value, used to estimate value from income. E.g. an office yielding 7% on €1,500/m^2 rent implies value ~€1,500/m^2 (because Rent/Value ≈ Yield). Lower yields mean higher value and usually reflect lower perceived risk.
- Absorption: The rate at which the market can "absorb" (sell or lease) the new space. We give in units/year or m^2/year. It's tied to demand. If we overshoot absorption, units stay unsold longer (which affects cash flow).
- Monte Carlo Simulation: A risk analysis technique where we randomly vary inputs within defined distributions many times to see a range (distribution) of outcomes (NPVs, etc.). It provides probabilities like "chance of NPV > 0".

Planning/Design Terms:
- GFA (Gross Floor Area): Total floor area of buildings (sum of all floors, measured to external walls), used for planning limits. We often give breakdown by use (residential GFA, etc.). It's different from net area (usable inside area). E.g., residential net might be ~80% of GFA if 20% is common areas[34].
- FAR (Floor Area Ratio or Index): The ratio of total GFA to land area. We have ~0.42 m^2/m^2 on average (46k/109k). Often given as index in local plans. Our site had subzone indices 0.4–0.7 as hints[97].
- ZUT (Zona Urbanistica di Trasformazione): A zoning designation in Turin for transformation areas – requiring special planning (like our convenzionato). It's essentially a redevelopment zone.
- Convenzionato: Short for Permesso di Costruire Convenzionato – a "contractual building permit". It's a planning tool in Italy where a developer signs an agreement with the city to provide certain public benefits/infrastructure in exchange for development rights. It's typical for large projects like this and is binding.
- PNRR: Italy's Piano Nazionale di Ripresa e Resilienza – funded by the EU Recovery Plan. It includes money for urban renewal (Torino has a program called "Torino Cambia" under PNRR). We mention it as a funding source.
- Housing Sociale: Refers to subsidized or affordable housing programs in Italy (usually moderate rental or sale housing with support, often co-funded by public entities).
- Anchor Tenant: A major, long-term tenant that provides stability (e.g. a supermarket or government office). Securing one often helps finance the project.
- Phase 1, 2, 3: Construction phases delineated for manageability and to match absorption. Phase 1 typically includes initial infrastructure because you need to unlock site access. Later phases add buildings sequentially.
- Innovation Hub/Campus: In our context, a cluster of buildings dedicated to research, education, startups, etc., likely partnering with academic or tech institutions. We use this term for Concept 3's unique component.
- Environment Park: A real entity next to our site – a science and technology park

focusing on environmental tech. We reference it as a synergy for offices or innovation uses.
- OMI Zone: Osservatorio del Mercato Immobiliare zones – Italian agency-defined areas for real estate stats. D9 is ours, used for price references.

Conventions & Units:
- All financial figures are in Euro (€). Millions are often shown as e.g. €34.5 M (which means €34,500,000).
- Areas: 1 ha (hectare) = 10,000 m^2. We sometimes say 5.7 ha of park = 57,000 m^2.
- Floors: When we say 5–7 floors, in Italian context that's roughly 5–7 stories above ground.
- Time: We reference years as Year 1, 2, ... relative to start of construction. If 2026 is start (Year 0 or 1), then Year 6 ~ 2031, etc. We assume today end of 2025 as Year 0 for NPV.
- Percentages: In use mix, % of GFA (e.g. 40% of GFA residential). In financial ratios, % margin or IRR as noted.
- Units vs. m^2: Residential program we gave in both m^2 and number of units. Where "~300 units" is stated, that is based on average size assumptions (e.g. 80–90 m^2 gross per unit in Concept 1)[34].
- Currency Year: We treated all costs and revenues in "2025 euros" with no inflation. In reality, costs might inflate but so might sale prices. This was a simplifying convention; effectively, it's a real-term analysis (with a nominal 11% discount that could be seen as 9% real + 2% inflation, roughly, since we had 2% inflation assumption but did not apply it). This is for feasibility stage consistency. Future analyses might switch to nominal cash flows with inflation explicitly modeled.

This concludes the report and its appendices. Every effort has been made to base conclusions on documented facts and clearly stated assumptions, following the Single Source of Truth principle. The recommendation is therefore well-grounded in the data at hand. The next project stage should refine these numbers with actual commitments and updated studies, at which point the decision can be revisited with higher confidence.

References

References (Chapters 2–3)

AACE International. (2020, August 7). 18R-97: Cost estimate classification system—As applied in engineering, procurement, and construction for the process industries (Recommended practice). AACE International.

D.M. 1444/1968. (1968). Decreto ministeriale 2 aprile 1968, n. 1444 (G.U. Serie Generale n. 97, 16-04-1968).

D.P.R. 380/2001. (2001). Decreto del Presidente della Repubblica 6 giugno 2001, n. 380: Testo unico delle disposizioni legislative e regolamentari in materia edilizia (G.U. Serie Generale n. 245, 20-10-2001; ripubblicato 15-11-2001).

International Valuation Standards Council. (2024). International valuation standards: Effective 31 January 2025.

Manganelli, B. (2015). Real estate investing: Market analysis, valuation techniques, and risk management. Springer. https://doi.org/10.1007/978-3-319-06397-3

Miles, M. E., Netherton, L. M., & Schmitz, A. (2015). Real estate development: Principles and process (5th ed.). Urban Land Institute.

Moorhead, M., Armitage, L., & Skitmore, M. (2024). Real estate development feasibility and hurdle rate selection. Buildings, 14(4), 1045. https://doi.org/10.3390/buildings14041045

Royal Institution of Chartered Surveyors. (2024). RICS valuation—Global standards (Red Book Global Standards).

Anthropic. (n.d.). Claude Opus. Retrieved December 2025, from https://www.anthropic.com/claude/opus

Archistar. (n.d.). Property development software. Retrieved December 2025, from https://www.archistar.ai/property-development-software/

Autodesk. (n.d.). Autodesk Forma: Free trial and education access information. Retrieved December 2025, from https://www.autodesk.com/products/forma/free-trial

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (pp. 610–623). Association for Computing Machinery. https://doi.org/10.1145/3442188.3445922

Boje, C., Guerriero, A., Kubicki, S., & Rezgui, Y. (2020). Towards a semantic construction digital twin: Directions for future research. Automation in Construction, 114, 103179. https://doi.org/10.1016/j.autcon.2020.103179

buildingSMART International. (n.d.). IFC (Industry Foundation Classes) standards and resources. Retrieved December 2025, from https://www.buildingsmart.org/standards/

CityBldr. (n.d.). CityBldr—We build smart cities. Retrieved December 2025, from https://www.citybldr.com/

El Jaouhari, A., Aouinti, N., & Oubbati, O. S. (2024). Automated valuation models: A systematic literature review of methods, applications, and challenges. International Journal of Strategic Property Management, 28(5), 286–301. https://doi.org/10.3846/ijspm.2024.22251

European Union. (2016). Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 (General Data Protection Regulation). Official Journal of the European Union.

European Union. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union.

Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep learning. MIT Press.

Google Cloud. (n.d.). Introducing Gemini Enterprise. Retrieved December 2025, from https://cloud.google.com/blog/products/ai-machine-learning/introducing-gemini-enterprise

International Organization for Standardization. (2018a). ISO 19650-1:2018—Organization and digitization of information about buildings and civil engineering works, including building information modelling—Information management using building information modelling—Part 1: Concepts and principles. ISO.

International Organization for Standardization. (2018b). ISO 19650-2:2018—Organization and digitization of information about buildings and civil engineering works, including building information modelling—Information management using building information modelling—Part 2: Delivery phase of the assets. ISO.

International Organization for Standardization. (2024). ISO 16739-1:2024—Industry Foundation Classes (IFC) for data sharing in the construction and facility management industries—Part 1: Data schema. ISO.

International Valuation Standards Council. (2024). International valuation standards: Effective 31 January 2025. IVSC.

Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, J., Madotto, A., & Fung, P. (2023). Survey of hallucination in natural language generation. ACM Computing Surveys, 55(12), Article 248. https://doi.org/10.1145/3571730

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In Advances in Neural Information Processing Systems 30 (pp. 4765–4774).

National Institute of Standards and Technology. (2023). Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1). https://doi.org/10.6028/NIST.AI.100-1

National Institute of Standards and Technology. (2024). Artificial intelligence risk management framework: Generative artificial intelligence profile (NIST AI 600-1). https://doi.org/10.6028/NIST.AI.600-1

OpenAI. (n.d.). ChatGPT release notes. Retrieved December 2025, from https://help.openai.com/en/articles/6825453-chatgpt-release-notes

OpenAI. (n.d.). File uploads FAQ. Retrieved December 2025, from https://help.openai.com/en/articles/8555545-file-uploads-faq

OpenAI. (n.d.). Models (OpenAI API documentation). Retrieved December 2025, from https://platform.openai.com/docs/models

OpenAI. (n.d.). Update to GPT-5 System Card: GPT-5.2 (PDF). Retrieved December 2025, from https://cdn.openai.com/pdf/3a4153c8-c748-4b71-8e31-aecbde944f8d/oai_5_2_system-card.pdf

Panko, R. R. (1998). What we know about spreadsheet errors. Journal of End User Computing, 10(2), 15–21. https://doi.org/10.4018/joeuc.1998040102

Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). “Why should I trust you?” Explaining the predictions of any classifier. In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 1135–1144). Association for Computing Machinery. https://doi.org/10.1145/2939672.2939778

Rosen, S. (1974). Hedonic prices and implicit markets: Product differentiation in pure competition. Journal of Political Economy, 82(1), 34–55. https://doi.org/10.1086/260169

Royal Institution of Chartered Surveyors. (2024). RICS valuation—Global standards (Red Book Global Standards). RICS.

Tagliaro, C., & Zanni, S. (2020). PropTech: Technology and innovation in the real estate industry. Journal of Property Investment & Finance, 38(4), 343–355. https://doi.org/10.1108/JPIF-12-2019-0068

TestFit. (n.d.). Role of AI in real estate development: Today and tomorrow. Retrieved December 2025, from https://www.testfit.io/blog/role-of-ai-in-real-estate-development

Ullah, F., Sepasgozar, S. M. E., Thaheem, M. J., & Al-Turjman, F. (2021). Barriers to the digitalisation and innovation of smart real estate: A managerial perspective on technology non-adoption. Environmental Technology & Innovation, 22, 101527. https://doi.org/10.1016/j.eti.2021.101527

References (Chapter 4)

European Union. (2016). Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation). Official Journal of the European Union.

European Union. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union.

International Organization for Standardization. (2018a). ISO 19650-1:2018. Organization and digitization of information about buildings and civil engineering works, including building information modelling - Information management using building information modelling - Part 1: Concepts and principles. ISO.

International Organization for Standardization. (2018b). ISO 19650-2:2018. Organization and digitization of information about buildings and civil engineering works, including building information modelling - Information management using building information modelling - Part 2: Delivery phase of the assets. ISO.

References (Chapter 6)

Altman, S. (2025). The Gentle Singularity. Blog post.

Amodei, D. (2024). Machines of Loving Grace: How AI Could Transform the World for the Better. Essay.

Bick, A., Blandin, A., & Deming, D. (2025). The State of Generative AI Adoption in 2025. Federal Reserve Bank of St. Louis, On the Economy (Nov 13, 2025).

McKinsey & Company. (2025). The State of AI: Global Survey 2025. McKinsey.

OpenAI. (2025). The State of Enterprise AI: 2025 Report. OpenAI.

Stanford University Institute for Human-Centered Artificial Intelligence. (2025). AI Index Report 2025. Stanford HAI.

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency.

European Union. (2016). Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 (General Data Protection Regulation).

European Union. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act).

International Organization for Standardization. (2018a). ISO 19650-1:2018. Organization and digitization of information about buildings and civil engineering works, including building information modelling - Information management using building information modelling - Part 1: Concepts and principles.

International Organization for Standardization. (2018b). ISO 19650-2:2018. Organization and digitization of information about buildings and civil engineering works, including building information modelling - Information management using building information modelling - Part 2: Delivery phase of the assets.

International Valuation Standards Council. (2024). International valuation standards: Effective 31 January 2025. IVSC.

Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, J., Madotto, A., & Fung, P. (2023). Survey of hallucination in natural language generation. ACM Computing Surveys, 55(12), 1-38.

National Institute of Standards and Technology. (2023). Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1).

Panko, R. R. (1998). What we know about spreadsheet errors. Journal of End User Computing, 10(2), 15-21.

Royal Institution of Chartered Surveyors. (2024). RICS valuation - Global standards (Red Book Global Standards). RICS.