

Article

The Cost of Uncoordinated Decisions in Early-Stage Architectural Design: A Game-Theoretic Proof of Concept

Alia Amer ^{1,*}  and Shady Attia ^{2,*} ¹ Faculty of Fine Arts, Capital University, Cairo 111568, Egypt² Sustainable Building Design Lab, Department of Urban and Environmental Engineering, Faculty of Applied Sciences, University of Liege, 4000 Liege, Belgium

* Correspondence: aliamer.hegazy@f-arts.capu.edu.eg (A.A.); shady.attia@uliege.be (S.A.)

Abstract

Early-stage architectural design distributes decisions across parties who control different variables but share the same performance consequences. Envelope geometry is settled largely on the architectural side, while glazing specification and system efficiency are settled largely on the engineering side, and the two sets of decisions determine thermal comfort, daylight, energy use, and cost jointly. Multi-objective optimization identifies which design states are attainable under a coordinated formulation, and multi-criteria methods rank alternatives under stated preferences, but neither represents what happens when coupled variables are held under separate utility functions. The objective of this paper is to test, as a proof of concept, whether modeling that division as a two-player non-cooperative game yields information that Pareto analysis of the same problem does not. A worked 3×3 game, built on constructed values consistent with documented physical relationships for a single-zone office, yields a unique pure strategy Nash equilibrium in which neither player holds a dominant strategy. All nine design states are Pareto-efficient in the four-objective performance space, but once the four indicators are aggregated into the two player utilities, the equilibrium is strictly Pareto-dominated in the two-player utility space: a different design state raises the first player's utility by 15.6 percent and the second player's utility by 22.0 percent simultaneously, a loss of 15.6 percent in the summed normalized utility of the two players relative to the best available joint state. This figure is a model-specific utility measure, not a share of energy, comfort, daylight, or cost performance. A sensitivity analysis across 169 weighting combinations shows the equilibrium migrating across eight distinct design states and falling outside the utility Pareto set in 54 of the 188 resulting equilibrium instances (29 percent), a descriptive frequency rather than a probability; 52 of the 169 combinations admit at least one such equilibrium. When the performance values and the preference weights are both perturbed at random, at least one inefficient equilibrium appears in 20 to 30 percent of draws, and a sequential Stackelberg version of the game reaches an inefficient state whichever player moves first. The results show that, within the specified utility representation, the stable state implied by decentralized decision-making can be inefficient and that, within a given model, the size of that gap can be estimated before the design is fixed. They do not show how large the gap is in real projects. The numerical demonstration uses constructed performance measures rather than simulation output and establishes the analytical behavior of the framework rather than the performance of any building. Applying the framework to simulated and, where possible, calibrated models is left to future work, and the framework is intended to inform design judgment, not to replace it.



Academic Editor: Majid Movahedi Rad

Received: 23 August 2026

Revised: 23 September 2026

Accepted: 24 September 2026

Published: 26 September 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons](https://creativecommons.org/licenses/by/4.0/)[Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: game theory; Nash equilibrium; Pareto efficiency; architectural design decision-making; decision support; multi-objective design; building performance; early-stage design; integrated design; design coordination

1. Introduction

Architectural design is a complex and decision-intensive process that involves multiple goals, interests, and sometimes conflicting priorities. Early-stage design decisions are essential, as they often influence running costs, functionality, aesthetics, and overall design performance. This is why simulation-supported decision methods have been pushed progressively earlier in the design timeline [1,2]. Napps et al. [3] describe the early phases as a pivotal juncture in the conception of a project, with choices made in this period wielding substantial influence over subsequent construction and operation, and the delivery arrangements governing how those decisions are taken have themselves become an object of study [4].

Computational support for these decisions rests on two families of methods. Multi-objective optimization searches a design space and returns a Pareto front of non-dominated solutions [5–7], and surrogate-assisted approaches have made that search tractable for realistic building models [8]. Multi-criteria decision-making aggregates competing objectives into a ranking through explicit preference weights [9,10], with the acknowledged difficulty that the resulting ranking changes with the weights [11]. Both families treat the design problem as a single optimization owned by one decision maker. Both are well established, and neither is in question here.

What neither represents is the structure of the decision itself. In practice, the variables that determine building performance are not held by one party. Li et al. [12] observe that under conventional design practice the cooperative operation between architects, structural engineers, and equipment engineers does not run smoothly, and that the professional division of labor leaves each designer responsible only for problems within their own discipline. Flager et al. [13] describe the same fragmentation from the tooling side, noting the long-standing inability of discipline-specific tools to share data effectively. Han et al. [14] add that professional barriers make direct algorithmic optimization difficult for designers to operate, so that architects seek assistance in making decisions rather than a computed final answer.

The two sets of decisions are nonetheless physically coupled. A glazing specification changes the comfort consequence of a window-to-wall ratio, and a window-to-wall ratio changes the energy return on a system upgrade. Each party optimizes what it controls against what it is measured on, and the resulting building is the joint product. A Pareto front computed over the union of all variables describes what a single coordinated decision maker could achieve. It does not describe the state implied by two parties optimizing separately.

That gap is what non-cooperative game theory addresses, and this paper asks a narrow, testable question: does equilibrium analysis of an architectural design problem produce information that Pareto analysis of the same problem does not? If the equilibrium necessarily coincides with the Pareto-efficient set of the same utility representation, the game-theoretic formulation adds little decision information and amounts to a relabeling. If the two diverge, the difference identifies a model-implied loss in utility that is associated with the represented division of decision authority and cannot be traced to any physical or budgetary constraint within the modeled feasible set. The objective of the paper is therefore methodological. It sets out the game formulation, applies it to a deliberately simple and

transparent example, and checks how far the result survives changes in its inputs. It does not attempt to measure the size of the coordination gap in real projects.

The paper makes four contributions. First, it formulates distributed early-stage architectural decisions as a two-player non-cooperative game in which the assignment of variables to players follows a common, though simplified, division between architectural and engineering decision domains. Second, it compares the resulting Nash equilibrium explicitly against both the four-objective Pareto set and the two-player utility Pareto set and quantifies the resulting loss in normalized utility. Third, it tests how equilibrium location and efficiency change with the preference weights and with random perturbations of the performance values. Fourth, it compares the simultaneous game with sequential Stackelberg variants in which either player moves first.

Section 2 reviews multi-objective optimization, multi-criteria methods, integrated design, and game-theoretic work in the building domain and locates the gap. Section 3 sets out the game formulation, the performance measures, the normalization procedure, and the analysis methods. Section 4 reports the payoff matrix, the equilibrium, the comparison with the four-objective and utility Pareto sets, the weight sensitivity, the Monte Carlo analysis, and the Stackelberg variants. Section 5 discusses what follows for design practice, including the place of the designer's judgment. Section 6 states the limits of the demonstration, and Section 7 concludes.

2. Related Work

This section reviews four bodies of work that intersect with the problem addressed here: multi-objective optimization of building performance, multi-criteria decision-making and the treatment of preference weights, integrated design and the distribution of decision authority across a project team, and game-theoretic approaches in construction and engineering design. The literature corpus assembled for this review, comprising the records screened and the basis on which each was retained, is openly deposited [15].

2.1. Multi-Objective Optimization in Building Design

Simulation-based optimization of building design is a mature field. Evins [5] surveyed 74 studies applying computational optimization to envelope design, system configuration, and renewable generation, while Machairas et al. [6] and Nguyen et al. [7] reviewed the algorithms and the coupling between the optimizer and simulation engine. Evolutionary methods, and NSGA-II in particular, dominate because the design problem is discrete, non-convex, and expensive to evaluate. Magnier and Haghighat [16] combined TRNSYS simulation with a genetic algorithm and a neural network surrogate to trade thermal comfort against energy use, and Basbagill et al. [17] applied multi-objective feedback to sequential conceptual design decisions.

The approach remains active and has grown more capable. Wu et al. [18] coupled a Bayesian-optimized gradient boosting surrogate with NSGA-II to generate a Pareto set over energy use, daylight, and thermal comfort for a residential case, and Li et al. [19] developed an early-stage framework optimizing energy, daylight, view quality, and thermal comfort across five Chinese climates. Cruz et al. [8] reviewed surrogate-assisted multi-objective building optimization systematically and drew a distinction that matters here. The weighted-sum approach aggregates normalized objectives into a single function and loses information about how the sub-objectives interfere, whereas the Pareto approach returns a range of non-dominated solutions and leaves the choice among them open. They also report that more than half of the studies surveyed stopped at the Pareto front, with fewer than half adopting any explicit process for selecting a preferred solution from it.

The Pareto front is the correct object when a single decision maker holds every variable and is free to select any point on it. Two lines of work qualify that assumption. Zhou et al. [20] point to a structural feature of the process itself: in a real design process, the variables are chosen sequentially, with those already fixed separated from those to be determined at later stages, so that a formulation treating all variables as simultaneously available misrepresents how the decision unfolds. Attia et al. [21] documented the obstacles to adoption in practice, including model uncertainty, computation time, and the learning curve. Østergård et al. [2] noted a further complication, namely that building design is a collaborative activity influenced by architects, engineers, contractors, and owners. The optimization literature registers this as a barrier to adoption. It does not model it.

2.2. Multi-Criteria Decision-Making and the Weighting Problem

Where objectives must be reduced to a single ranking, multi-criteria methods assign explicit weights and utilizes aggregation. [9]. Villalba et al. [22] reviewed 91 building-related applications and found the analytic hierarchy process, simple additive weighting, and TOPSIS to be the prevalent instruments across assessment, selection, and retrofit. Grierson [10] combined optimization with multi-criteria ranking and treated the weight vector as the point at which stakeholder priorities enter the calculation.

The standing objection is that a result that moves with the weights carries limited authority unless the movement itself is reported, and the building literature now documents that movement directly. Rachman et al. [11] compared five multi-criteria methods for passive design selection and tested rank stability under systematic variation of the criterion weights, finding the analytic hierarchy process comparatively robust to weight variability and VIKOR the most sensitive. Theilig et al. [23] reached a compatible conclusion for the renovation of exterior wall structures, reporting that utility analysis and TOPSIS are more sensitive to emphasized criteria weighting than the alternatives tested. Han et al. [14] approach the same problem from the designer's side, proposing an interactive method through which architects modify a scheme according to their own preferences rather than accepting an algorithmically determined result.

This paper adopts the discipline that follows from these findings. Section 4.5 treats weight sensitivity as a result to be reported rather than a caveat to be acknowledged, because if published comparisons of established methods find their rankings to be weight-dependent, a new framework has no basis for presenting a single weighting as authoritative.

2.3. Integrated Design and Distributed Decision Authority

The framework proposed here assigns design variables to two players. That partition is a simplification, but it is grounded in an established observation about how building design work is organized, and the integrated design literature documents both the division of responsibility and its consequences for performance.

Napps et al. [3] characterize the early design phase as requiring the collaborative integration of architecture, structural engineering, energy analysis, and cost planning, and evaluate design variants across spatial, structural, and environmental criteria precisely because no single discipline holds all of the relevant information. Li et al. [12] address what follows when that integration is absent. Reviewing integrated design for climate-responsive buildings, they observe that the cooperative operation between architects, structural engineers, and equipment engineers does not run smoothly under conventional practice, with consequences for whole-building performance. Their analysis of the constraints on integrated design identifies the clear professional division of labor as a primary obstacle, noting that each designer is responsible only for problems within their own discipline and that, in the absence of conflict, there is limited information exchange between professions.

The consequences are documented rather than hypothetical. Flager et al. [13] report survey evidence that architects and engineers take more than a month to complete a single design cycle and average fewer than three cycles during conceptual design, with most engineers using simulation to validate a chosen alternative rather than to explore several. Attia et al. [24] find that these barriers persist, reporting that architects and engineers seeking design decision aid experience decisional conflict, and that cognitive, technological, and social factors continue to limit the uptake of simulation-based support in practice. They formalize multidisciplinary design optimization as an attempt to structure problem decomposition and coordination among groups working on a complex engineering system, which requires identifying the actors responsible for each step and the information passed between them. Mollaoglu-Korkmaz et al. [4] provide empirical support from twelve case studies, linking project delivery method to the level of team integration achieved and, in turn, to sustainability outcomes.

This is the situation the present study formalizes. Envelope geometry is settled largely on the architectural side, while specification and system efficiency are settled largely on the engineering side, and each side is accountable for a different set of performance measures. The physical consequences do not follow that division. What the integrated design literature describes as a coordination problem corresponds, in game-theoretic terms, to coupled decisions held under separate utility functions. Section 3 states that structure formally, while noting in Section 3.1 that the resulting players remain analytical constructs rather than models of individual behavior.

2.4. Game-Theoretic Approaches in Construction and Building Design

Game theory has an established presence in construction engineering and management. Eissa et al. [25] mapped the field and found applications concentrated in bidding, contract selection, risk allocation, and dispute resolution, that is, in situations where the players are contracting parties with divergent financial interests. Yoon and Pishdad-Bozorgi [26] used a game-theoretic framework to analyze tacit knowledge sharing and procurement choice in mega projects.

A second strand addresses stakeholder behavior in building energy markets. Zheng et al. [27] constructed a tripartite evolutionary game among governments, energy service companies, and clients to examine energy performance contracting in public buildings, and identified clients' overestimation of potential loss as the critical barrier to adoption. Gao et al. [28] combined prospect theory with evolutionary game theory for green renovation of existing residential buildings and reported that the evolutionary outcomes are highly sensitive to the prospect theory parameters, with high loss aversion driving the system toward a suboptimal equilibrium. Their finding that equilibrium location depends materially on preference parameters is directly relevant to Section 4.5 of this paper, although their players are market actors rather than design disciplines.

A third strand operates on the building itself rather than on the market around it. Tang and Wang [29] modeled demand-side flexibility engagement as a Stackelberg game between an aggregator and building users, with transactional service prices and flexibility capacities determined at Nash equilibrium and users' dissatisfaction costs quantified through comfort indicators. This is the closest existing formulation to the present one, in that it couples comfort against economics across parties with distinct objectives inside a building performance problem. It differs in that it concerns operational engagement in a completed building rather than design decisions, and in that its players are an aggregator and its customers rather than design disciplines.

A fourth strand applies equilibrium concepts to engineering design rather than to contracting or operation. Marousi and Charitopoulos [30] surveyed this territory for pro-

cess and energy systems engineering, observing that the field has traditionally emphasized centralized decision-making schemes while decentralized schemes have attracted increasing attention for their ability to capture multi-stakeholder dynamics more accurately, and reviewing the Nash, Stackelberg, cooperative, and bilevel formulations through which that shift has been implemented. Within engineering design specifically, Tang et al. [31] formulated multi-criterion aerodynamic shape design as a Nash game between design variable subsets, demonstrating that splitting variables between players and solving for equilibrium gives a tractable route through a high-dimensional coupled problem, and Martínez-Muñoz et al. [32] applied game-theoretic multi-objective optimization to environmental and social life cycle assessment of composite bridges.

2.5. Research Gap and Positioning

These four studies answer four different questions. Optimization asks which design states are attainable [5–8,16–20]. Multi-criteria analysis asks which alternatives are preferred under stated preferences, and how far that ranking survives changes in those preferences [9–11,22,23]. Integrated design research asks how decision authority is distributed across a project team and what the distribution costs [3,4,12–14]. Game theory asks which states are stable when coupled decisions are held under separate utility functions [25–32].

Game theory is therefore already established in the building domain, in three areas distinct from the present one. One body of work models stakeholder behavior in energy retrofit and renovation markets, where the players are governments, energy service companies, and clients, and the object of analysis is contract adoption [27,28]. A second models operational interaction between building users and aggregators, trading comfort against cost at runtime in a completed building [29]. A third, in process and engineering design, treats variable partitions as players in order to make coupled optimization tractable [30–32]. Across this literature, we did not identify work that partitions coupled early-stage architectural and engineering design variables according to professional decision authority and then compares the resulting Nash equilibrium directly against the Pareto-efficient set in order to quantify a coordination-related efficiency gap. That comparison is the contribution offered here.

One boundary deserves to be stated explicitly. Li et al. [19] employed a hybrid weighting method to accommodate diverse project demands and report that objective priorities differ by climate and facade orientation, which resembles the weight variation reported in Section 4.5 of this paper. The two exercises use the same lever for different purposes. Their variation tailors a recommended design to project context. The variation here tests whether the equilibrium remains stable and whether it remains efficient, which is a question about the decision structure rather than about the design.

This study is accordingly a methodological contribution to architectural decision support, not a theoretical contribution to game theory. It introduces no new solution concept and makes no claim to supersede optimization or multi-criteria methods, which answer different questions and answer them well. The framework is complementary: optimization identifies what is attainable, and equilibrium analysis identifies what is stable given who decides what. The gap between those two quantities is the object of interest.

3. Materials and Methods

The core methodological workflow is summarized in Figure 1. It proceeds from defining the design agents, strategy sets, and performance objectives to constructing the payoff matrix and identifying the Nash equilibrium. The subsequent subsections extend

this workflow by evaluating Pareto efficiency, preference weight sensitivity, robustness to performance perturbations, and sequential Stackelberg outcomes.




Figure 1. Methodological framework for applying game theory and Nash equilibrium to architectural design decision-making.

3.1. Game Formulation

The design process is modeled as a two-player non-cooperative game in strategic form, $G = \{A, B, S_A, S_B, U_A, U_B\}$, where A and B are the players, S_A and S_B are their finite strategy sets, and U_A and U_B are their utility functions. A joint strategy $x = (x_A, x_B)$ with $x_A \in S_A$ and $x_B \in S_B$ defines one complete design state. The game is non-cooperative because the two players hold different utility functions and enter no binding agreement to coordinate: each selects the strategy that maximizes its own utility given the strategy of the other. Standard treatments of the solution concepts used here are given by Osborne and Rubinstein [33] and Fudenberg and Tirole [34].

A joint strategy (x_A^*, x_B^*) is a pure strategy Nash equilibrium [35] when $U_A(x_A^*, x_B^*) \geq U_A(x_A, x_B^*)$ for all $x_A \in S_A$ and $U_B(x_A^*, x_B^*) \geq U_B(x_A^*, x_B)$ for all $x_B \in S_B$. Neither player can raise its own utility by changing strategy alone. This is a condition of stability under unilateral deviation and does not imply joint efficiency, a distinction that Section 4.4 makes central.

The two players are analytical agents that stand for two contrasting design logics. They are not models of an individual architect or engineer. Three notions are kept apart in what follows. Variable ownership is a modeling choice about which agent sets which

variable. Professional responsibility is an institutional matter, and in practice architects, engineers, clients, and cost consultants share responsibility for comfort, daylight, energy, and cost. Utility is the set of objectives on which an agent is scored. The partition used here draws on the observation that envelope geometry tends to be settled on the architectural side and specification on the engineering side [12,13]. That literature describes a tendency, not a fixed rule, and no data on how responsibility for these particular criteria is divided in real teams were collected for this study. The equilibrium is therefore read as the outcome of two design logics pursued separately, which may or may not match how a given team is organized. The equilibrium remains a property of the formulated model and its utility functions and is not a behavioral prediction about any individual.

3.2. Players, Decision Domains, and Strategies

Player A is the passive environmental design agent and controls facade geometry, represented in this demonstration by the window-to-wall ratio and the depth of a horizontal shading device. Player B is the energy and resource design agent and controls the specification of the envelope and the systems, represented by the glazing thermal and optical properties (U-value, solar heat gain coefficient, and visible light transmittance), the external wall U-value, and the seasonal efficiency of the heating and cooling system. These are the variables that change across the strategies in Table 1. Other variables that belong to the same decision domains in practice, such as orientation and thermal mass on the architectural side and airtightness on the engineering side, are held constant across all nine joint strategies and are not varied by either player in this demonstration. Each variable is assigned to exactly one player, which fixes strategic ownership. The physical effects of a variable are not confined to its owner: the window-to-wall ratio is controlled by A, but changes energy use and cost, which enter B's utility, and glazing specification is controlled by B but changes comfort and daylight, which enter A's utility. That asymmetry between control and consequence is the source of strategic interdependence and is what makes the problem a game rather than two independent optimizations.

Table 1. Player strategies and controlled design variables.


Player/Strategy	Description	Controlled Variables
A1	Compact glazing, deep shading	WWR 20%, horizontal shading depth 0.9 m
A2	Moderate glazing, moderate shading	WWR 35%, horizontal shading depth 0.6 m
A3	Extensive glazing, light shading	WWR 55%, horizontal shading depth 0.3 m
B1	Baseline envelope and systems	Clear double glazing (U 2.7 W/m ² K, SHGC 0.70, VLT 0.78); wall U 0.45 W/m ² K; HVAC SCOP 2.6
B2	Solar control envelope, efficient systems	Solar control double glazing (U 1.4 W/m ² K, SHGC 0.35, VLT 0.62); wall U 0.24 W/m ² K; HVAC SCOP 3.4
B3	High-performance envelope and systems	Triple glazing, high selectivity (U 0.8 W/m ² K, SHGC 0.30, VLT 0.68); wall U 0.15 W/m ² K; heat pump SCOP 4.2

SHGC, solar heat gain coefficient; VLT, visible light transmittance; SCOP, seasonal coefficient of performance. Thermal mass, and airtightness are held constant across all strategies.

Three strategies per player are defined for the demonstration, listed in Table 1. They span the range an early-stage design team would realistically consider for a single-zone office and are deliberately coarse, consistent with the resolution available before the concept is fixed. Figure 2 shows the three facade options and the three envelope and system specifications schematically. Each Player B strategy bundles glazing, wall, and system upgrades, so B1 to B3 represent three coherent specification levels rather than independent choices of each component. System efficiency is expressed for all three levels as the seasonal coefficient

cient of performance (SCOP), so that the three values are directly comparable. The values are illustrative of standard, improved, and high-efficiency heat generation equipment.

(a) Player A: facade geometry (elevation and section)



(b) Player B: envelope and system specification (glazing and wall sections)




Figure 2. Schematic of the design strategies. (a) Facade elevations and sections for Player A strategies A1 to A3, showing window-to-wall ratio and horizontal shading depth. (b) Glazing, wall, and system specifications for Player B strategies B1 to B3. The drawings are not to scale.

3.3. Performance Objectives and Metrics

Four objectives are evaluated for each joint strategy. Thermal comfort PT is the proportion of occupied hours with predicted mean vote inside the acceptable band. Daylight PD is spatial daylight autonomy, the proportion of floor area meeting the target illuminance for the required fraction of occupied hours, evaluated with operable blinds so that glare-driven blind deployment reduces the score. Energy PE is annual energy use intensity. Cost PC is the initial construction cost of the envelope and systems. In an applied study, these would be produced by dynamic simulation, for instance EnergyPlus (26.1.0) for thermal and energy performance [36] and Radiance-based methods for daylight [37], with cost from a quantity take-off. Their treatment in this paper is set out in Section 3.7.

Player A holds utility over thermal comfort and daylight. Player B holds utility over energy and cost. This split is a stylized reading of how accountability is often divided in a project team (Section 3.1), not a claim that each profession cares only about these criteria. It is this split in utility, not any physical decoupling, that generates the divergence reported in Section 4.4.

3.4. Normalization

The four objectives carry different units and opposite directions of preference, so each is normalized to the unit interval by min–max scaling across the nine joint strategies. For objectives where a higher physical value is preferable, Equation (1) applies. For objectives where a lower physical value is preferable, which is the case for energy use and cost, Equation (2) applies.

$$P_i = (x_i - \min(x)) / (\max(x) - \min(x)) \quad (1)$$

$$P_i = (\max(x) - x_i) / (\max(x) - \min(x)) \quad (2)$$

Min–max scaling is one of several available procedures. Vafaei et al. [38] note that criteria measured in different units must be transformed into comparable dimensionless values, commonly on the unit interval, and that the choice among techniques such as max, max–min, sum, and vector normalization can itself alter the resulting ranking. It is adopted here for transparency and because it makes the consistency check described below available, not because it is uniquely correct. After normalization, $0 \leq P_i \leq 1$ for every objective, and a higher value always denotes better performance. Min–max scaling requires that each normalized column attain both 0 and 1 across the evaluated set, which provides an internal consistency check on the performance table and is verified in Section 4.1. Normalization is applied to performance, not to preference: a normalized performance measure states how good a design state is on one objective, whereas a utility states how much a given player values that state. The two are distinct and are kept distinct throughout.

3.5. Utility Functions

Each player’s utility is a weighted sum of the two normalized measures in its domain:

$$UA(XA, XB) = \omega_{TPT}(XA, XB) + \omega_{D PD}(XA, XB) \quad (3)$$

$$UB(XA, XB) = \omega_{EPE}(XA, XB) + \omega_{C PC}(XA, XB) \quad (4)$$

with $\omega_T + \omega_D = 1$ and $\omega_E + \omega_C = 1$. The base case sets all four weights to 0.5, which is a neutral reference rather than a claim that the objectives matter equally. Section 4.5 varies them systematically and reports how far the conclusions depend on the choice.

3.6. Equilibrium and Efficiency Analysis

Best responses are read column-wise for Player A and row-wise for Player B. A strategy is dominant when it is a best response to every strategy of the opponent; where a dominant strategy exists, the game is solvable by iterated elimination, and equilibrium analysis contributes nothing beyond it, so the absence of dominance is verified before the equilibrium is reported. Cells that are mutual best responses are pure strategy Nash equilibria. Where none exists, a mixed-strategy equilibrium in probability distributions over pure strategies is guaranteed for finite games; its interpretation in a design context is discussed in Section 5.4.

Efficiency is assessed separately. A joint strategy is Pareto-efficient in the two-player utility space when no other joint strategy raises one player’s utility without lowering the other’s. Where the equilibrium lies outside this utility Pareto set, the welfare loss is quantified as the ratio of the maximum attainable sum of utilities to the sum at equilibrium, following the price of anarchy measure introduced by Koutsoupias and Papadimitriou [39].

Two Pareto sets are distinguished. The four-objective Pareto set contains the joint strategies that are non-dominated in the four-objective performance space, that is, on the four normalized measures PT, PD, PE, and PC. The utility Pareto set contains the joint strategies that are non-dominated on the two utilities UA and UB. With positive weights, the utility Pareto set is always a subset of the four-objective Pareto set, because a state dominated on all four measures is also dominated on both utilities. The reverse does not hold: aggregation within each player hides the trade-off between comfort and daylight and the trade-off between energy and cost. Both sets are reported in Section 4.4 so that what the aggregation adds, and what it hides, can be seen directly. In what follows, statements that an equilibrium is Pareto-dominated, inefficient, or outside the Pareto set refer to the two-player utility space unless the four-objective performance space is named explicitly.

The sum $UA + UB$ is used as the welfare measure because it is the usual basis of the price of anarchy [39,40] and because both utilities lie on the same normalized scale. The sum carries value judgments. It assumes that a unit of utility for one player is worth a unit for the other, and it inherits the min–max normalization and the weights. The resulting percentage is therefore a loss in normalized utility within this model. It is not a percentage of energy use, comfort hours, daylight area, or cost. Where an equilibrium is strictly Pareto-dominated, its inefficiency does not depend on the summation, since any welfare measure that increases in both utilities ranks the dominating state higher. The size of the loss does depend on the measure, and Section 4.4 reports it under the sum, the product, and the minimum of the two utilities.

3.7. Status of the Numerical Values

The performance measures in Table 2 were constructed rather than simulated. They were built to satisfy the physical relations that govern this problem, namely that comfort improves with shading and with envelope quality, that the comfort penalty of a large glazed area shrinks as the envelope improves, an interaction Li et al. [19] report from optimization across five climates when they find that increased window-to-wall ratio becomes beneficial in colder conditions because high-performance envelope settings mitigate the associated heat transfer, that daylight rises with glazed area but is degraded by blind deployment under high-transmittance glazing and by the reduced visible transmittance of solar control coatings, that energy savings from system upgrades diminish at the margin while capital costs escalate, and that the return on a system upgrade is larger for a heavily glazed facade. They also satisfy the min–max consistency condition of Section 3.4. Table A1 in Appendix A reports illustrative physical values for thermal comfort, daylight, annual energy use intensity, and initial construction cost that correspond to the normalized entries of Table 2. They were obtained by applying Equations (1) and (2) in reverse over a stated reference range for each objective, so that the relationships described here can be checked in physical units. The daylight relation likewise reflects documented shading behavior rather than an assumed one. O'Brien et al. [41] report that most office occupants adjust shades no more than weekly or monthly, responding to long-term solar trends rather than to short-term events, and cite measured evidence that blind position can alter the useful contribution of daylight by between 10 and 60 percentage points. Van Den Wymelenberg [42] reaches a compatible conclusion from an independent review, observing that decisions about blind position are formed over weeks or months rather than hours, and that substantially less daylight enters when blinds are down, so that daylight-sensing lighting controls save correspondingly less energy. The constructed values accordingly assign lower useful daylight performance to heavily glazed cases paired with high-transmittance glazing, where glare exposure is greatest, and shades are most likely to remain closed. The relations are therefore constructed rather than measured, but they are not arbitrary, and each of the two principal mechanisms has independent support in the literature. What follows establishes the analytical behavior of the framework under physically coherent inputs. It does not report the performance of any building, and the specific equilibrium location is a property of these values rather than a finding about offices in general. Section 6 states what a simulated case study would add. The paper should therefore be read as a methodological proof of concept. The numbers show that a coordination gap can arise under physically coherent inputs. They do not show how often it arises, or how large it is, in real projects.

Table 2. Normalized performance measures for the nine joint strategies.

Joint Strategy	Thermal Comfort P_T	Daylight P_D	Energy P_E	Cost P_C
(A1, B1)	0.72	0.10	0.30	0.82
(A1, B2)	0.88	0.00	0.58	0.52
(A1, B3)	1.00	0.08	0.86	0.10
(A2, B1)	0.40	0.52	0.34	1.00
(A2, B2)	0.70	0.44	0.66	0.72
(A2, B3)	0.92	0.56	1.00	0.22
(A3, B1)	0.00	0.62	0.00	0.86
(A3, B2)	0.48	0.80	0.46	0.54
(A3, B3)	0.80	1.00	0.86	0.00

Higher is better for all four objectives after normalization. Corresponding illustrative physical values are given in Table A1 (Appendix A).

3.8. Robustness Analyses

Two further analyses test whether the divergence depends on the specific inputs. The first is a Monte Carlo analysis of the performance values. In each of 10,000 draws, every entry in Table 2 receives an independent perturbation drawn uniformly from the interval $\pm\delta$; negative values are set to zero, and each column is normalized again by min–max scaling so that the consistency condition of Section 3.4 still holds. Three ranges are tested, $\delta = 0.05, 0.10$, and 0.20 , which correspond to 5, 10, and 20 percent of the range of each objective across the nine states. Each range is run twice, once with equal weights and once with ω_T and ω_E drawn uniformly between 0.20 and 0.80 in every draw. For each draw, the pure strategy equilibria and the utility Pareto set are recomputed, and a draw is counted as inefficient when at least one of its pure strategy equilibria lie outside the utility Pareto set. The ranges are illustrative. They were not derived from simulation uncertainty, which would require the simulated case study discussed in Section 6, and independent perturbations ignore the physical correlations between objectives that a simulation would preserve.

The second analysis replaces simultaneous moves with sequential ones. In a Stackelberg game, the leader chooses a strategy while anticipating the follower’s best response, and the follower then responds to the leader’s choice [30]. Two orders are tested with the base-case payoffs. When Player A leads, the facade concept is fixed before the specification. When Player B leads, the specification or the budget for it is set first.

4. Results

4.1. Normalized Performance Measures

Table 2 gives the normalized performance of the nine joint strategies on the four objectives. Every column reaches both 0 and 1, satisfying the min–max condition. The interaction that drives the rest of the analysis is visible in the thermal comfort column: the penalty for moving from A1 to A3 is 0.72 under B1 but only 0.16 under B3, because a high-performance envelope absorbs the consequence of a larger glazed area. The daylight column shows the countervailing effect, with A3 scoring 0.62 under B1 rather than its maximum, because clear high-transmittance glazing on a heavily glazed facade drives blind deployment and suppresses useful daylight.

4.2. Payoff Matrix

Applying Equations (3) and (4) with equal weights gives the payoff matrix in Table 3. Each cell holds the ordered pair (UA, UB). For example, (A1, B1) has PT = 0.72 and PD = 0.10, so $UA = 0.5(0.72) + 0.5(0.10) = 0.41$, and PE = 0.30 and PC = 0.82, so $UB = 0.5(0.30) + 0.5(0.82) = 0.56$.

Table 3. Two-player payoff matrix under equal weights. Each cell gives the ordered pair (UA, UB).

	B1	B2	B3
A1	(0.41, 0.56)	(0.44, 0.55)	(0.54, 0.48)
A2	(0.46, 0.67)	(0.57, 0.69)	(0.74, 0.61)
A3	(0.31, 0.43)	(0.64, 0.50)	(0.90, 0.43)

Rows are Player A strategies and columns are Player B strategies; UA and UB are computed with Equations (3) and (4) and all weights equal to 0.5.

4.3. Best Responses and Equilibrium

Reading Table 3 column-wise, Player A's best response is A2 against B1 (0.46 against 0.41 and 0.31) and A3 against both B2 (0.64) and B3 (0.90). Reading row-wise, Player B's best response is B1 against A1 (0.56) and B2 against both A2 (0.69) and A3 (0.50). Neither player holds a dominant strategy: A2 is optimal against B1 but not against B2 or B3, and B1 is optimal against A1 but not against A2 or A3. The game cannot be reduced by iterated elimination and requires equilibrium analysis.

The single cell that is a mutual best response is (A3, B2), with payoffs (0.640, 0.500). This is the unique pure strategy Nash equilibrium. The three near misses are informative. At (A2, B1), Player A is playing its best response, but Player B would move to B2. At (A3, B3), Player A is content, but Player B would retreat to B2 to avoid the capital cost of triple glazing. At (A1, B1), Player B is content, but Player A would move to A2 for daylight. In each case, one party is satisfied, and the other is not, which is the structure that drives the design toward the equilibrium.

4.4. Equilibrium Against the Utility Pareto Set

The utility Pareto set computed over the nine states contains three joint strategies: (A2, B2) at (0.57, 0.69), (A2, B3) at (0.74, 0.61), and (A3, B3) at (0.90, 0.43). The equilibrium (A3, B2) at (0.640, 0.500) is not among them. It is strictly dominated in the two-player utility space by (A2, B3), which delivers 0.740 to Player A and 0.610 to Player B, improving both simultaneously by 15.6 percent and 22.0 percent, respectively.

The sum of utilities at equilibrium is 1.140 against a maximum of 1.350 at (A2, B3), a price of anarchy of 1.184, which corresponds to a loss of 15.6 percent in summed normalized utility. Measured by the product of the two utilities, the loss is 29.1 percent (0.320 against 0.451), and measured by the lower of the two utilities, it is 18.0 percent (0.500 against 0.610). None of these values is a percentage of energy, comfort, daylight, or cost performance. The loss can be traced to the design variables. At equilibrium, Player A has taken the largest glazed area (A3) while Player B has stopped at the mid-specification envelope (B2). Neither can improve alone: reducing glazing costs Player A daylight, and upgrading to triple glazing costs Player B capital without a compensating energy return at that glazing area. The dominating state (A2, B3) requires both to move at once, with Player A accepting a smaller glazed area and Player B accepting a higher specification, at which point the reduced solar load and the improved envelope combine to deliver more of both objectives. No single party can propose that move and benefit from it; it is available only to a coordinated decision.

This is the result the framework exists to produce. A Pareto analysis of the same nine states identifies (A2, B3) as attainable but is silent on whether it will be reached. Equilibrium analysis shows that it is not the stable state implied by the modeled decision structure and puts a number on the shortfall in utility terms, conditional on the utility functions, weights, and performance values used.

The picture changes when the four objectives are examined before aggregation. On PT, PD, PE, and PC, none of the nine joint strategies is dominated, because no state is matched or beaten on all four measures by another state. The four-objective Pareto set therefore contains all nine states, including the equilibrium, so the equilibrium is not inefficient with respect to the four building-performance objectives themselves. The dominance of (A2, B3) over (A3, B2) exists only after aggregation. (A2, B3) is better on thermal comfort (0.92 against 0.48) and energy (1.00 against 0.46), whereas (A3, B2) is better on daylight (0.80 against 0.56) and cost (0.54 against 0.22). Under equal weights, each player gains more on one of its objectives than it loses on the other, so both utilities rise. Two consequences follow. First, for this problem, the four-objective front does not discriminate between states, since it excludes none of them. The utility Pareto set narrows the choice to three states, and the equilibrium analysis indicates which state the modeled decision structure settles on. Second, the inefficiency is conditional on the weights inside each utility. A player who valued daylight far more than comfort would not regard (A2, B3) as an improvement, which is the situation examined in Section 4.5. Figure 3 shows the nine states in utility space.




Figure 3. The nine joint strategies in utility space under equal weights. The utility Pareto set is joined by a line. The Nash equilibrium (A3, B2), which is also the Stackelberg outcome when Player A leads, and the Stackelberg outcome when Player B leads, (A2, B1), both lie below it. The dashed arrow shows the joint move to (A2, B3), which raises both utilities. All nine states belong to the four-objective Pareto set.

4.5. Sensitivity to Preference Weights

The weights are the most contestable input, so the game was re-solved across a grid of ω_T and ω_E from 0.20 to 0.80 in steps of 0.05, giving 169 weighting combinations. Every combination admitted at least one pure strategy Nash equilibrium, so no mixed-strategy analysis was required.

The equilibrium location is highly sensitive to the weights. Eight of the nine design states appear as an equilibrium somewhere in the grid, and the results are summarized in Table 4. The equilibrium sat at (A3, B3) in 36 percent of cases and at (A2, B1) in 28 percent, and no single state accounts for a majority. Because some weighting combinations admit more than one equilibrium, the 169 combinations yield 188 equilibrium instances in total. The equilibrium lay outside the utility Pareto set in 54 of these 188 instances, that is, 29 percent. This is an instance-level descriptive frequency: a combination with two equilibria contributes two observations, and the figure should not be read as a probability of inefficiency. Counted by weighting combination, 19 combinations admit two equilibria, and none admit more than two; 52 of the 169 combinations (30.8 percent) admit at least one equilibrium outside the utility Pareto set, and in 39 (23.1 percent) every equilibrium lies outside it. The states most vulnerable to inefficiency are (A1, B1), which was inefficient in 30 of 35 occurrences, and (A3, B2), the base-case equilibrium, which was inefficient in 15 of 27.

Table 4. Equilibrium location and efficiency across 169 weighting combinations.

Equilibrium State	Occurrences	Share of Grid	Pareto-Efficient
(A3, B3)	60	36%	60 of 60
(A2, B1)	48	28%	48 of 48
(A1, B1)	35	21%	5 of 35
(A3, B2)	27	16%	12 of 27
(A3, B1)	5	3%	0 of 5
(A2, B3)	5	3%	5 of 5
(A2, B2)	4	2%	4 of 4
(A1, B2)	4	2%	0 of 4

Occurrences count equilibrium instances (188 in total). Shares are expressed per weighting combination and sum to more than 100 percent because 19 combinations admit two equilibria. Pareto-efficient refers to the two-player utility Pareto set.

Two conclusions follow. The specific equilibrium reported in Section 4.3 is a property of the equal-weight assumption and should not be read as a general result about office design. The structural finding holds across the tested grid: for weights between 0.20 and 0.80, the equilibrium of the modeled two-party process lies off the utility Pareto set in 29 percent of equilibrium instances, and at least one such equilibrium occurs in 31 percent of weighting combinations. The framework locates the weighting regions in which that risk is highest. The range from 0.20 to 0.80 is a numerical test range, chosen to avoid cases in which one objective is ignored altogether. It is not an empirically established range of designer preferences. For a project team, the practical output is not the equilibrium itself but the identification of the regions where the model-implied loss from uncoordinated decisions is largest.

4.6. Robustness to Perturbed Performance Values

Table 5 summarizes the Monte Carlo analysis. With equal weights and small perturbations ($\delta = 0.05$), the base-case equilibrium (A3, B2) survives in 97.6 percent of draws, and at least one equilibrium lies off the utility Pareto set in 97.7 percent of draws. As the perturbation grows, the payoff structure changes more often. At $\delta = 0.20$, the base-case equilibrium survives in 37.6 percent of draws and an inefficient equilibrium appears in 31.9 percent. The high share at $\delta = 0.05$ with equal weights reflects the persistence of the inefficient base-case equilibrium, and the share falls as larger perturbations displace

it. When the weights are drawn at random as well, the share of draws with at least one inefficient equilibrium lies between 20.3 and 30.4 percent across the three ranges. This is of the same order as the 31 percent of weighting combinations with at least one inefficient equilibrium on the weight grid, although the two figures come from different sampling schemes and are not directly comparable. The range of 20 to 30 percent therefore applies to joint perturbation of values and weights, not to perturbation of the performance values alone. The share of draws with no pure strategy equilibrium rises from 0.5 to 13.5 percent with equal weights as δ increases. Where an inefficient equilibrium occurs, the median loss in summed normalized utility lies between 15.8 and 21.0 percent.

Table 5. Monte Carlo analysis of perturbed performance values (10,000 draws per row).

Range δ	Weights	No Pure Equilibrium	Base-Case Equilibrium Persists	Draws with at Least One Inefficient Equilibrium	Median Loss (P10 to P90)
0.05	Equal	0.5%	97.6%	97.7%	16.4% (12.9 to 20.3)
0.05	Random	1.2%	18.2%	30.4%	15.8% (9.0 to 22.8)
0.10	Equal	6.5%	66.6%	65.1%	17.3% (11.0 to 25.2)
0.10	Random	4.0%	16.7%	25.6%	17.4% (9.8 to 24.8)
0.20	Equal	13.5%	37.6%	31.9%	21.0% (11.0 to 32.8)
0.20	Random	9.1%	18.0%	20.3%	20.4% (10.5 to 31.4)

An inefficient draw has at least one pure strategy equilibrium outside the utility Pareto set. Losses are in summed normalized utility and are computed only for inefficient draws. Equal weights: all four weights set to 0.5. Random weights: ω_T and ω_E drawn uniformly between 0.20 and 0.80.

Two readings follow. The base-case equilibrium is fragile once the inputs move by more than a few percent of their range, so the location of the equilibrium should not be taken from constructed values. The divergence itself is not the product of one carefully tuned table: under joint perturbation of values and weights, between one draw in three and one draw in five produces an equilibrium that both players could improve on together. Whether that share is representative of real design problems cannot be established from constructed data.

4.7. Sequential Moves: Stackelberg Variants

With Player A as leader, A anticipates from Table 3 that B will answer A1 with B1, A2 with B2, and A3 with B2. Its utilities from these three outcomes are 0.41, 0.57, and 0.64, so it chooses A3, and the outcome is (A3, B2), the same state as the simultaneous equilibrium, with a summed utility of 1.14. With Player B as leader, B anticipates that A will respond to B1 with A2, B2 with A3, and B3 with A3. Its utilities from these outcomes are 0.67, 0.50, and 0.43, so it chooses B1, and the outcome is (A2, B1), with utilities (0.46, 0.67) and a sum of 1.13. That state also lies off the utility Pareto set, because (A2, B2) gives (0.57, 0.69) and improves both players. In this example, moving from simultaneous to sequential decisions changes which inefficient state is reached but does not remove the inefficiency. The order of moves decides who carries most of the loss: Player B obtains 0.50 when A leads, and Player A obtains 0.46 when B leads. A one-shot leader and follower is still a simplification of a design process that iterates and revisits earlier choices.

5. Discussion

5.1. What Equilibrium Adds Beyond Pareto Analysis

The question posed in Section 1 has an answer for this case. The equilibrium lies outside the two-player utility Pareto set, although it belongs to the four-objective Pareto set,

so the two analyses do not coincide, and the difference between them is interpretable. The utility Pareto set describes what the modeled strategies could deliver to both players under a coordinated choice. The equilibrium identifies the stable state implied by decentralized decision-making under the stated utility functions and decision structure. The gap between the two is a model-implied coordination inefficiency. It cannot be attributed to a physical or budgetary constraint, because both states lie inside the same modeled feasible set, but its existence and size depend on the specified utility functions, strategy set, normalization procedure, preference weights, and constructed performance relationships. The analysis does not establish that professional fragmentation alone causes a gap of the observed magnitude. The measure used to express it is standard in the analysis of distributed engineered systems. Chandan et al. [40] define the price of anarchy as the loss in efficiency caused by distributed decision-making, computed as the ratio between social welfare at the system optimum and at equilibrium, and treat the misalignment between agents' local utilities and the system objective as its source. Perakis and Roels [43] apply the same measure to decentralized supply chains, quantifying how far a chain of autonomous parties falls short of the fully coordinated equivalent. The situation in Section 4.4 has the same formal structure, with the architectural and engineering design logics in place of the contracting parties and normalized performance utility in place of profit.

This distinction matters because the two analyses recommend different actions. A Pareto analysis returns a set of good design states and leaves the selection to preference. An equilibrium analysis returns the state that is stable under unilateral deviation and identifies whether a coordinated intervention is worth making. Where the equilibrium is already efficient, as it was in 71 percent of the equilibrium instances tested, the analysis finds no joint move among the modeled strategies that would raise both utilities. That is a narrow statement. It does not mean that the design meets its comfort, energy, or cost requirements, that the strategy set contains the best available options, or that coordination on other variables and with other parties is unnecessary. Where the equilibrium is inefficient, the analysis names the joint move that would raise both utilities and estimates its value in utility terms, which gives a project team a specific item to discuss.

5.2. Implications for Integrated Design

Integrated design processes are usually argued for on general grounds. The analysis here supplies a candidate mechanism and a way of estimating its size within a model. In the base case, the loss is 15.6 percent of summed normalized utility. It arises because each player makes the choice that is best for its own objectives, not because either makes an error. An architectural logic that keeps the larger glazed area and an engineering logic that declines the triple glazing are both consistent within their own objectives. Within the model, the inefficiency arises from the assumed division of decision authority in combination with the specified utilities, weights, and performance values, and more effort inside either domain does not remove it.

That reframes what an integrated design workshop is for. Its function is not to improve either party's judgment but to make available the joint moves that neither can initiate alone. That framing is consistent with evidence linking the level of team integration achieved on a project to its sustainability outcomes [4], and with the finding that architects seek methods that support their decisions rather than replacing them [14]. The framework provides a way to identify those moves before the workshop takes place, and to establish whether the exercise is warranted on this project at all. A shared performance target or a combined budget can change the problem in two distinct ways. If the two utilities are replaced by a single aggregated utility, the game becomes a centralized single-objective optimization whose solution is one state that depends on the chosen aggregation and weights; it is not a

Pareto front. If instead the four performance objectives are kept separate under a single decision maker, the relevant object is the four-objective Pareto set, which in this example contains all nine states and leaves the choice to preferences applied afterward.

5.3. Relation to Optimization and Multi-Criteria Methods

In conventional optimization approaches, architects identify the best possible outcome by maximizing or minimizing a specific objective. When multiple objectives are involved, these are typically combined into a single function through the assignment of weights. The outcome then depends on the weighting scheme, and changes in the weights lead to different design solutions [9,10]. The same dependence applies to the framework presented here, which is why Section 4.5 reports it rather than assuming it away. Comparative studies of established multi-criteria methods in the building domain find exactly this behavior, with rank order shifting under systematic weight variation and the methods differing in how far they resist it [11,23].

The relationship between the methods is complementary and sequential rather than competitive. Optimization is the appropriate tool for generating the design states and their performance; the framework presented here takes those states as input and asks a question about the decision structure that optimization does not address. In an applied workflow, a simulation-based optimization of the kind reviewed by Cruz et al. [8] would produce the performance table, the game-theoretic layer would identify the equilibrium and the efficiency gap, and multi-criteria methods would remain available for selecting among efficient states once coordination is established. This division of labor is consistent with the observation that most optimization studies stop at the Pareto front without an explicit procedure for choosing among its members [8].

5.4. Mixed Strategies

Every weighting combination tested admitted a pure strategy equilibrium, so mixed strategies did not arise. They are retained in the formulation because finite games guarantee a mixed-strategy equilibrium where no pure one exists, but their interpretation in design requires care. A mixed strategy is a probability distribution over strategies, not a physical blend of design alternatives: a 0.4 weight on A1 and 0.6 on A2 does not describe a window-to-wall ratio between the two. It describes uncertainty about which strategy a party will adopt. Where such an equilibrium arises, the useful reading is that the design decision is genuinely unresolved given the information available, which is itself a finding worth reporting to a client.

5.5. The Place of the Designer

A framework of this kind raises a fair question: can architectural design be reduced to a calculation, and if it could, what would be left of the designer's work and sensibility? Our position is that it cannot, and the framework does not attempt to do so. Rittel and Webber [44] argued that planning and design problems cannot be fully stated in advance and have no single correct solution. Schön [45] described design as a reflective conversation with the situation, in which the designer frames the problem while working on it, and Lawson [46] showed how, in architecture, problem and solution develop together. The technical and the humanistic sides of architecture cannot be separated in practice. The size of a window is at the same time a question of heat gain, daylight, view, proportion, privacy, and the character of the facade.

The game-theoretic model works on a small part of that problem once it has already been framed. It takes as given the options that designers have proposed, the objectives that can be measured, and the weights that someone has chosen. What makes a proposal good architecture sits outside the model: how the options were generated, spatial and formal

quality, meaning, and the relation to place and to the people who use the building. The contribution is narrow. For the measurable objectives, the model shows whether separately made decisions could leave both sides worse off than a joint decision, and which joint move would help. That result is an input to design discussion. It can bring the architect and the engineer to the table earlier than they might otherwise meet, but whether a smaller glazed area is acceptable, and what it does to the building, remains a matter of design judgment.

6. Limitations

The performance measures are constructed rather than simulated. They are physically coherent and internally consistent, and they are sufficient to establish that the equilibrium and the utility Pareto set can diverge and to measure the divergence under stated assumptions. They are not evidence about any real building, and the equilibrium location reported in Section 4.3 should not be transferred to other cases. Establishing the magnitude of the coordination gap in practice requires a simulated case study with EnergyPlus and a Radiance-based daylight model on a defined zone, climate, and cost basis, with the uncertainty in such model inputs treated explicitly [47]. That is the immediate next step for this work. The Monte Carlo analysis in Section 4.6 tests sensitivity to the values but cannot substitute for such a study, because independent random perturbations do not respect the physical correlations a simulation would preserve. A simulation study could also start from a model calibrated against measured data. Calibration cannot be applied to a building that does not yet exist, but a calibrated model of a comparable existing building could serve as the reference from which the design options are varied, which would reduce the risk that the coordination gap is an artifact of a poorly chosen baseline. Calibration has its own limits: accepted criteria mostly address energy use rather than comfort or daylight, several different models can meet the criteria for the same measured data [48], and errors in applying the statistical validation indices can give the false impression that a model is calibrated [49]. Calibration against monitored indoor conditions has nevertheless been used to predict thermal comfort across a housing stock [50] and to compare envelope retrofit options for comfort and energy [51], which shows how such a reference model could be extended beyond energy use. The remaining calibration uncertainty would still need to be carried through to the equilibrium analysis.

Three further limits apply. The utility functions are linear and additive within each domain, which excludes threshold effects such as a comfort criterion that must be met rather than traded. The main analysis treats the game as simultaneous and one-shot, whereas design usually proceeds in sequence. Zhou et al. [20] documented this directly, noting that design variables are chosen progressively across stages, with those already fixed separated from those still to be determined. Section 4.7 tests the simplest sequential alternative, a Stackelberg game with either player as leader, of the kind surveyed by Marousi and Charitopoulos [30] and applied to building flexibility by Tang and Wang [29]. Real projects iterate, revisit earlier decisions, and share information imperfectly, which a single leader and follower do not capture. Feng et al. [52] provided a construction domain precedent, modeling a construction supply chain as a Stackelberg game between hierarchical decision makers and showing that a nominally centralized optimum can be unstable because a lower-level actor is free to deviate in its own interest. A further limitation follows from the shading assumption. The constructed daylight values presume manually operated shades. Automated and integrated daylighting controls behave differently and can improve energy performance and visual comfort relative to manual operation [53], so the daylight penalty assigned here to heavily glazed cases would be smaller in a building with automated shading. The partition into two players is a simplification of a team that in practice includes the client, the contractor, and the facility operator, each with distinct utilities; extending

to three or more players is straightforward in principle, but the payoff structure becomes harder to present and to interpret. The assignment of objectives to players is also stylized, and the results would change if, for example, cost entered both utilities.

7. Conclusions

This paper formulated early-stage architectural design as a two-player non-cooperative game in which the assignment of variables to players follows a simplified division between architectural and engineering decision logics, and asked whether equilibrium analysis yields information that Pareto analysis does not.

In this example, it does. In the worked case, neither player holds a dominant strategy, the game has a unique pure strategy Nash equilibrium, and that equilibrium, although it belongs to the four-objective Pareto set, is strictly Pareto-dominated in the two-player utility space by a design state that improves both players simultaneously, by 15.6 percent and 22.0 percent, for a loss of 15.6 percent in summed normalized utility. The dominating state is reachable only by a simultaneous move that neither party can initiate unilaterally and from which neither benefits by moving alone. Across 169 weighting combinations, the equilibrium migrated across eight design states and fell outside the utility Pareto set in 29 percent of the resulting equilibrium instances (54 of 188), and 52 of the 169 combinations admitted at least one such equilibrium. When the performance values and the weights were both perturbed at random, at least one inefficient equilibrium appeared in 20 to 30 percent of draws, and a Stackelberg version of the game reached an inefficient state whichever player moved first. These results describe a model-implied inefficiency under the stated representation of decision authority and preferences, not a measured effect of professional fragmentation.

The practical implication is modest. Part of the gap between what a design team could achieve and what it settles on may come from the way decisions are divided, and a model of this kind can flag where that risk exists, and which joint moves would address it. The size of that gap in real projects has not been measured here, and the framework is meant to support design judgment, not to replace it.

The demonstration rests on constructed rather than simulated performance data and establishes the analytical behavior of the framework rather than the performance of any building. Applying it to a simulated, and ideally calibrated, single-zone office across several climates and extending the formulation to sequential and multi-party settings are the next steps.

Supplementary Materials: The following supporting information can be downloaded at: <https://www.mdpi.com/article/10.3390/buildings16193823/s1>.

Author Contributions: Conceptualization, A.A. and S.A.; methodology, A.A. and S.A.; formal analysis, A.A.; investigation, A.A.; writing, original draft preparation, A.A.; writing, review and editing, S.A.; visualization, A.A.; supervision, S.A. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: All numerical values supporting the reported results are contained in Tables 2–5 and in Table A1 (Appendix A). The script used to construct the performance table, solve the game, and run the weight sensitivity, Monte Carlo (random seed 2026), and Stackelberg analyses is available from the corresponding author on request. The literature corpus underlying Section 2 is openly available in the Harvard Dataverse at <https://doi.org/10.7910/DVN/ALZQER>

(accessed on 27 August 2026), reference [15]. The Python (3.14.7) verification script used to perform and verify the numerical game-theoretic analysis is provided as Supplementary Material (game_analysis_verification.py), enabling independent verification of the reported calculations and results.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Illustrative Physical Values of the Constructed Performance Matrix

Table A1 lists illustrative physical values that correspond to the normalized performance measures in Table 2. Each column was obtained by applying Equation (1) or (2) in reverse over a stated reference range, so that normalizing the columns of Table A1 again with Equations (1) and (2) returns Table 2 exactly. The values allow the relationships described in Section 3.7 to be read in physical units. Moving from B1 to B3 reduces energy use intensity by 44.8 kWh/m²·yr under A1 but by 68.8 kWh/m²·yr under A3, so the return on the specification upgrade is larger for the heavily glazed facade. The thermal comfort penalty of moving from A1 to A3 falls from 28.8 percentage points under B1 to 8.0 percentage points under B3. The values are constructed for the demonstration and are not simulation or cost plan outputs.

Table A1. Illustrative physical values corresponding to the normalized performance measures in Table 2.

Joint Strategy	Thermal Comfort (% of Occupied Hours)	Daylight, sDA (% of Floor Area)	Energy Use Intensity (kWh/m ² ·yr)	Envelope and System Cost (€/m ²)
(A1, B1)	83.8	35.0	126.0	296
(A1, B2)	90.2	30.0	103.6	356
(A1, B3)	95.0	34.0	81.2	440
(A2, B1)	71.0	56.0	122.8	260
(A2, B2)	83.0	52.0	97.2	316
(A2, B3)	91.8	58.0	70.0	416
(A3, B1)	55.0	61.0	150.0	288
(A3, B2)	74.2	70.0	113.2	352
(A3, B3)	87.0	80.0	81.2	460

Reference ranges used to invert the min–max normalization: thermal comfort 55 to 95 percent of occupied hours within the acceptable PMV band; daylight 30 to 80 percent of floor area (spatial daylight autonomy); energy use intensity 70 to 150 kWh/m²·yr; initial construction cost of envelope and systems 260 to 460 €/m² of floor area. Lower values are preferable for energy use intensity and cost.

References

1. Attia, S.; Gratia, E.; De Herde, A.; Hensen, J.L.M. Simulation-based decision support tool for early stages of zero-energy building design. *Energy Build.* **2012**, *49*, 2–15. [\[CrossRef\]](#)
2. Østergård, T.; Jensen, R.L.; Maagaard, S.E. Building simulations supporting decision making in early design: A review. *Renew. Sustain. Energy Rev.* **2016**, *61*, 187–201. [\[CrossRef\]](#)
3. Napps, D.; Staudt, J.; Saluz, U.; Chen, X.; Steiner, D.; Zong, C.; Deghim, F.; Lang, W.; Geyer, P.; Schnellenbach-Held, M.; et al. Evaluation of building design variants in early phases on the basis of adaptive detailing strategies. *Buildings* **2026**, *16*, 685. [\[CrossRef\]](#)
4. Mollaoglu-Korkmaz, S.; Swarup, L.; Riley, D. Delivering sustainable, high-performance buildings: Influence of project delivery methods on integration and project outcomes. *J. Manag. Eng.* **2013**, *29*, 71–78. [\[CrossRef\]](#)
5. Evins, R. A review of computational optimisation methods applied to sustainable building design. *Renew. Sustain. Energy Rev.* **2013**, *22*, 230–245. [\[CrossRef\]](#)
6. Machairas, V.; Tsangrassoulis, A.; Axarli, K. Algorithms for optimization of building design: A review. *Renew. Sustain. Energy Rev.* **2014**, *31*, 101–112. [\[CrossRef\]](#)

7. Nguyen, A.-T.; Reiter, S.; Rigo, P. A review on simulation-based optimization methods applied to building performance analysis. *Appl. Energy* **2014**, *113*, 1043–1058. [\[CrossRef\]](#)
8. Cruz, A.S.; Caldas, L.R.; Mendes, V.M.; Mendes, J.C.; Bastos, L.E.G. Multi-objective optimization based on surrogate models for sustainable building design: A systematic literature review. *Build. Environ.* **2024**, *266*, 112147. [\[CrossRef\]](#)
9. Marler, R.T.; Arora, J.S. Survey of multi-objective optimization methods for engineering. *Struct. Multidiscip. Optim.* **2004**, *26*, 369–395. [\[CrossRef\]](#)
10. Grierson, D.E. Pareto multi-criteria decision making. *Adv. Eng. Inform.* **2008**, *22*, 371–384. [\[CrossRef\]](#)
11. Rachman, A.P.; Ichwania, C.; Mangkuto, R.A.; Pradipta, J.; Koerniawan, M.D.; Sarwono, J. Comparison of multi-criteria decision-making methods for selection of optimum passive design strategy. *Energy Build.* **2024**, *314*, 114285. [\[CrossRef\]](#)
12. Li, Z.; Tian, M.; Zhu, X.; Xie, S.; He, X. A review of integrated design process for building climate responsiveness. *Energies* **2022**, *15*, 7133. [\[CrossRef\]](#)
13. Flager, F.; Welle, B.; Bansal, P.; Soremekun, G.; Haymaker, J. Multidisciplinary process integration and design optimization of a classroom building. *J. Inf. Technol. Constr.* **2009**, *14*, 595–612.
14. Han, Z.; Li, X.; Sun, J.; Wang, M.; Liu, G. An interactive multi-criteria decision-making method for building performance design. *Energy Build.* **2023**, *282*, 112793. [\[CrossRef\]](#)
15. Amer Hegazy, A.; Attia, S. *Dataset of Literature for Game-Theoretic Coordination in Building Performance Design, V2*; Harvard Dataverse; Harvard University: Cambridge, MA, USA, 2026. [\[CrossRef\]](#)
16. Magnier, L.; Haghighat, F. Multiobjective optimization of building design using TRNSYS simulations, genetic algorithm, and artificial neural network. *Build. Environ.* **2010**, *45*, 739–746. [\[CrossRef\]](#)
17. Basbagill, J.P.; Flager, F.L.; Lepech, M. A multi-objective feedback approach for evaluating sequential conceptual building design decisions. *Autom. Constr.* **2014**, *45*, 136–150. [\[CrossRef\]](#)
18. Wu, C.; Pan, H.; Luo, Z.; Liu, C.; Huang, H. Multi-objective optimization of residential building energy consumption, daylighting, and thermal comfort based on BO-XGBoost-NSGA-II. *Build. Environ.* **2024**, *254*, 111386. [\[CrossRef\]](#)
19. Li, L.; Qi, Z.; Ma, Q.; Gao, W.; Wei, X. Evolving multi-objective optimization framework for early-stage building design: Improving energy efficiency, daylighting, view quality, and thermal comfort. *Build. Simul.* **2024**, *17*, 2097–2123. [\[CrossRef\]](#)
20. Zhou, Y.; Tam, V.W.Y.; Le, K.N. Developing a multi-objective optimization model for improving building's environmental performance over the whole design process. *Build. Environ.* **2023**, *246*, 110996. [\[CrossRef\]](#)
21. Attia, S.; Hamdy, M.; O'Brien, W.; Carlucci, S. Assessing gaps and needs for integrating building performance optimization tools in net zero energy buildings design. *Energy Build.* **2013**, *60*, 110–124. [\[CrossRef\]](#)
22. Villalba, P.; Sánchez-Garrido, A.J.; Yepes, V. A review of multi-criteria decision-making methods for building assessment, selection, and retrofit. *J. Civ. Eng. Manag.* **2024**, *30*, 465–480. [\[CrossRef\]](#)
23. Theilig, K.; Vollmer, M.; Lang, W.; Albus, J. Multi-criteria decision-making for energy building renovation: Comparing exterior wall structures with the AHP, ANP, utility analysis, and TOPSIS. *Build. Environ.* **2025**, *280*, 113075. [\[CrossRef\]](#)
24. Attia, S.; El-Degwy, A.; Attia, M. Usability and fitness testing for building performance simulation tools. *J. Build. Perform. Simul.* **2024**. [\[CrossRef\]](#)
25. Eissa, R.; Eid, M.S.; Elbeltagi, E. Current applications of game theory in construction engineering and management research: A social network analysis approach. *J. Constr. Eng. Manag.* **2021**, *147*, 04021066. [\[CrossRef\]](#)
26. Yoon, J.H.; Pishdad-Bozorgi, P. Game theory-based framework for analyzing the collaborative dynamic of tacit knowledge sharing and the choice of procurement and contract types in mega construction projects. *Buildings* **2022**, *12*, 305. [\[CrossRef\]](#)
27. Zheng, S.; Zhou, Y.; Yuan, J.; Liu, R.; Lyu, P.; Han, Z.; Zhang, C. Understanding governments, ESCOs, and clients' behavioral strategies in public building energy-efficiency renovation based on evolutionary game theory. *J. Manag. Eng.* **2025**, *41*, 04025018. [\[CrossRef\]](#)
28. Gao, Y.; Liu, J.; Zhang, J.; Xie, H. Modeling the evolutionary mechanism of multi-stakeholder decision-making in the green renovation of existing residential buildings in China. *Buildings* **2025**, *15*, 2758. [\[CrossRef\]](#)
29. Tang, H.; Wang, S. Game-theoretic optimization of demand-side flexibility engagement considering the perspectives of different stakeholders and multiple flexibility services. *Appl. Energy* **2023**, *332*, 120550. [\[CrossRef\]](#)
30. Marousi, A.; Charitopoulos, V.M. Game theoretic optimization in process and energy systems engineering: A review. *Front. Chem. Eng.* **2023**, *5*, 1130568. [\[CrossRef\]](#)
31. Tang, Z.; Désidéri, J.-A.; Périaux, J. Multicriterion aerodynamic shape design optimization and inverse problems using control theory and Nash games. *J. Optim. Theory Appl.* **2007**, *135*, 599–622. [\[CrossRef\]](#)
32. Martínez-Muñoz, D.; Martí, J.V.; Yepes, V. Game-theory-based multi-objective optimization for enhancing environmental and social life cycle assessment in steel-concrete composite bridges. *Mathematics* **2025**, *13*, 273. [\[CrossRef\]](#)
33. Osborne, M.J.; Rubinstein, A. *A Course in Game Theory*; MIT Press: Cambridge, MA, USA, 1994.
34. Fudenberg, D.; Tirole, J. *Game Theory*; MIT Press: Cambridge, MA, USA, 1991.
35. Nash, J. Non-cooperative games. *Ann. Math.* **1951**, *54*, 286–295. [\[CrossRef\]](#)

36. Crawley, D.B.; Lawrie, L.K.; Winkelmann, F.C.; Buhl, W.F.; Huang, Y.J.; Pedersen, C.O.; Strand, R.K.; Liesen, R.J.; Fisher, D.E.; Witte, M.J.; et al. EnergyPlus: Creating a new generation building energy simulation program. *Energy Build.* **2001**, *33*, 319–331. [[CrossRef](#)]
37. Reinhart, C.F.; Walkenhorst, O. Validation of dynamic RADIANCE-based daylight simulations for a test office with external blinds. *Energy Build.* **2001**, *33*, 683–697. [[CrossRef](#)]
38. Vafaei, N.; Ribeiro, R.A.; Camarinha-Matos, L.M. Assessing normalization techniques for simple additive weighting method. *Procedia Comput. Sci.* **2022**, *199*, 1229–1236. [[CrossRef](#)]
39. Koutsoupias, E.; Papadimitriou, C. Worst-case equilibria. *Comput. Sci. Rev.* **2009**, *3*, 65–69. [[CrossRef](#)]
40. Chandan, R.; Paccagnan, D.; Marden, J.R. Methodologies for quantifying and optimizing the price of anarchy. *IEEE Trans. Autom. Control* **2024**, *69*, 7742–7757. [[CrossRef](#)]
41. O'Brien, W.; Kapsis, K.; Athienitis, A.K. Manually-operated window shade patterns in office buildings: A critical review. *Build. Environ.* **2013**, *60*, 319–338. [[CrossRef](#)]
42. Van Den Wymelenberg, K. Patterns of occupant interaction with window blinds: A literature review. *Energy Build.* **2012**, *51*, 165–176. [[CrossRef](#)]
43. Perakis, G.; Roels, G. The price of anarchy in supply chains: Quantifying the efficiency of price-only contracts. *Manag. Sci.* **2007**, *53*, 1249–1268. [[CrossRef](#)]
44. Rittel, H.W.J.; Webber, M.M. Dilemmas in a general theory of planning. *Policy Sci.* **1973**, *4*, 155–169. [[CrossRef](#)]
45. Schön, D.A. *The Reflective Practitioner: How Professionals Think in Action*; Basic Books: New York, NY, USA, 1983.
46. Lawson, B. *How Designers Think: The Design Process Demystified*, 4th ed.; Architectural Press: Oxford, UK, 2005.
47. Hopfe, C.J.; Hensen, J.L.M. Uncertainty analysis in building performance simulation for design support. *Energy Build.* **2011**, *43*, 2798–2805. [[CrossRef](#)]
48. Coakley, D.; Raftery, P.; Keane, M. A review of methods to match building energy simulation models to measured data. *Renew. Sustain. Energy Rev.* **2014**, *37*, 123–141. [[CrossRef](#)]
49. Ruiz, G.R.; Bandera, C.F. Validation of calibrated energy models: Common errors. *Energies* **2017**, *10*, 1587. [[CrossRef](#)]
50. Calama-González, C.M.; Suárez, R.; León-Rodríguez, Á.L. Thermal comfort prediction of the existing housing stock in southern Spain through calibrated and validated parameterized simulation models. *Energy Build.* **2022**, *254*, 111562. [[CrossRef](#)]
51. López-Escamilla, Á.; Herrera-Limones, R.; León-Rodríguez, Á.L. Double-skin facades for thermal comfort and energy efficiency in Mediterranean climate buildings: Rehabilitating vulnerable neighbourhoods. *Buildings* **2024**, *14*, 326. [[CrossRef](#)]
52. Feng, C.; Ma, Y.; Zhou, G.; Ni, T. Stackelberg game optimization for integrated production-distribution-construction system in construction supply chain. *Knowl.-Based Syst.* **2018**, *157*, 52–67. [[CrossRef](#)]
53. Konstantoglou, M.; Tsangrassoulis, A. Dynamic operation of daylighting and shading systems: A literature review. *Renew. Sustain. Energy Rev.* **2016**, *60*, 268–283. [[CrossRef](#)]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.