Airport Design Arbitrations

Aviation Volatility and Contractual Permanence

Dr Karim Hafez[*]

I. An Airport Is Not a Building 2
II. The Hidden Airline: Airport Design as Business Model 4
A. Fleet Mix and the False Simplicity of Capacity 5
B. Lounges, Hierarchy, and the Passenger the Airport Imagines 6
III. Procurement: Who Owns the Assumption? 7
A. Design-Bid-Build and the Refuge of Execution 7
B. Design-and-Build and the Illusion of Single-Point Certainty 8
C. Management Contracting and the Vanishing Integrator 8
IV. The Airport Beyond the Terminal 9
A. ATC, Slots, and the External Limits of Capacity 9
B. Fuel, Fire, Rescue, and the Hidden Choreography of the Apron 10
C. Pavement, Climate, and the Legal Life of Materials 11
V. Dwell, Minutes, and the Opening That Is Not an Opening 12
A. The Commercial Airport and the Paradox of Dwell 12
B. Minimum Connection Time and the Tyranny of Minutes 13
C. Completion, Certification, and the Opening That Is Not an Opening 14
VI. Proof: Systems Failure, Models, and the Paper Airport 15
A. The Failure of Linear Causation 15
B. Modelling Evidence and the Tyranny of Assumptions 16
C. ORAT: When the Paper Airport Meets the Real Airport 17
VII. The Expandable Airport: Designing for a Future That Must Arrive on Time 18
A. The Capacity Figure and the Illusion of Precision 18
B. Granular Mandates and the Problem of Specified Futures 19
C. The Expanding Airport and the Presumptive Airline Problem, Revisited 20

---

[*] The author holds an FAA Commercial Pilot Licence with multiple ratings.

D. Structural Expandability, Systemic Expandability, and the Gap Between Them 21
E. The Masterplan as Moving Target 22
F. Advocacy and Determination in Expansion Cases 23
VIII. Law: Standard of Care, Applicable Law, and the Frontier of Loss 24
A. Standard of Care and the Operational Undertaking 24
B. Contract, Applicable Law, and Ex-Contractual Routes 25
C. Strategic Economic Loss and the Identity of the Employer 26
IX. From Airport to Award: The Tribunal's Task 27
X. Conclusion: The Impossible Stability of Airports 29

I. An Airport Is Not a Building

An airport is not a building: it is a wager on the future of aviation dressed in concrete and glass.

That distinction is not rhetorical: a building is static in ambition, designed to exist, and existence is enough; an airport is kinetic in ambition, designed to perform, and performance depends upon an aviation world that will not stand still. The building need only survive: the airport must synchronise, and synchronisation, unlike survival, can fail in a thousand ways simultaneously.

Consider what airport success actually requires: the time between arrival and departure; the time a bag moves from one aircraft hold to another; the time a stand remains occupied after pushback; the time lost at transfer security; the time a fuel truck blocks a service road; the time inserted by ATC between two departures on the same or intersecting runways; the time required by rescue services to reach a runway end; the time a passenger may be made to dwell profitably without incurring a missed connection. None of these is a structural property: all of them are temporal relationships between moving parts. A building has no equivalent: the airport is, at its core, a machine for managing time, and time, unlike steel and glass, cannot as easily be fixed in a contract.

Most construction disputes concern structures; airport design disputes concern assumptions: the forecasts, models, operating concepts and commercial propositions that the structure was designed to serve. The bridge carries weight; the weight is calculable; the bridge either bears it or fails. The airport carries uncertainty: traffic that may or may not materialise; airlines that may migrate, collapse or grow beyond all expectation; security conditions that may transform overnight; climates that shift across the life of a pavement; passenger behaviours that no model fully anticipates. These forces are not peripheral to the airport's purpose: they are its operating

environment, and to design an airport is to make predictions about all of them — simultaneously.

This article advances a single argument, from which all else follows: airport design arbitrations are distinctively difficult because they arise at the point where aviation volatility collides with contractual permanence. The contract and the masterplan aspire to fix, for years or decades, arrangements that aviation reality will continue to revise: arbitration begins when that aspiration fails — or, more precisely, when the parties disagree about who bore the risk that it would.

That collision — volatility meeting permanence — is not the backdrop: it is the constitutive feature of airport arbitration, the reason why capacity is rarely just capacity; why procurement structures often conceal rather than settle responsibility; why proof so often becomes proof of systems behaviour rather than individual defect; why the standard of care grows unstable when operational success is at stake; and why alleged losses strain the border between recoverable project damage and diffuse public economic consequence. Each of these difficulties is a distinct expression of the same underlying tension; none can be understood in isolation from it.

A word on the scope of what follows. The observations in this article have been developed over two decades and across several airport disputes. They do not reflect any single project, jurisdiction or arbitration, and they should not be read as doing so: every airport case is configuration-sensitive, and the configurations vary enormously. Procurement structure, the actual terms of the contract, the applicable law, airline mix, security posture, climatic setting, ATC environment, operating concept and the identity of the employer may each alter the analysis substantially. What follows is a framework, not a formula: a map of the fault-lines, not a guarantee that any particular fault will move in any particular direction. The same symptom may have different causes; the same cause may produce different legal consequences; and the analysis appropriate to one airport may be actively misleading when transplanted to another.

The argument proceeds in stages: the airport reconceived not as a structure but as a synchronised system; the hidden airline inside the masterplan; procurement as an attempted and always imperfect allocation of operational assumption; the airport's external constraints — ATC, slots, fuel, rescue, pavement and climate; the paradox of dwell and the tyranny of the minute; proof — the difficulty of linear causation, modelling evidence and ORAT; the legal framework — standard of care, the means-result distinction, applicable law and the frontier of strategic economic loss; and, finally, the tribunal's task stated with the precision it demands.

The order is deliberate: the article does not begin with law because the legal problem cannot be accurately stated until the airport has been properly understood. If the airport is treated as a building, the dispute will seem familiar: a variant of delay or defect, amenable to the usual grammar of construction arbitration. If it is treated as a synchronised aviation system, the familiar categories begin to strain: delay ceases to be only lateness and becomes lost sequencing; defect ceases to be only non-compliance

and becomes impaired operation; capacity ceases to be a number and becomes a temporal and behavioural relation between fixed design and moving aviation reality. That reclassification is the article's hinge: without it, the rest is detail; with it, the detail begins to acquire analytical force.

II. The Hidden Airline: Airport Design as Business Model

Every airport masterplan contains an identified or presumptive airline: not as metaphor, but as a description of how airports are actually designed.

Terminal geometry is not neutral: it assumes a carrier — or at least a type of carrier — and arranges its corridors, piers, lounges, security lanes, baggage halls and retail zones accordingly. A low-cost, point-to-point airport imagines speed: short turns, high aircraft utilisation, austerity in amenity, minimal transfer complexity and restrained capital expenditure; the passenger is a unit of throughput; the building is a processing machine; efficiency is measured in minutes between landing and boarding the next aircraft. A long-haul hub imagines something altogether different: banked arrivals generating transfer waves; premium lounges absorbing passengers who will wait an hour and expect to find it tolerable; alliance-aware security and immigration channels; through-baggage systems capable of making connections that the passenger cannot; redundancy against irregular operations because one disrupted bank can corrupt several hours of schedule. A cargo-oriented airport imagines yet another logic: airside road access, warehouse integration, customs adjacency, night operations, load-bearing apron surfaces and the irrelevance of most of what the passenger terminal regards as essential. A tourist gateway imagines still another: seasonal surge, charter traffic, bus stands at the perimeter, immigration queues that the summer leisure market will endure even when the business traveller would not.

The airport presents itself as neutral infrastructure; it is not: its every dimension has already taken sides. The pier width, the stand mix, the lounge square footage, the security lane count, the baggage belt configuration, the corridor width at the transfer point — all of these are decisions, and every decision encodes a prediction about the airline and passenger world the airport is designed to serve: design is strategy in physical form, and strategy, when aviation reality overtakes it, becomes the subject matter of arbitration.

The legal significance of this is simple but frequently overlooked: when the airport fails to perform as designed, the first question is not whether the contractor built it correctly, but whether the design's embedded airline prediction was ever the contractor's responsibility to validate, challenge or guarantee. That question depends upon the contract and the law; it cannot even be asked until the prediction has been identified.

A. Fleet Mix and the False Simplicity of Capacity

Capacity is the most treacherous word in airport disputes: it appears objective and is not.

Ten million annual passengers are not ten million annual passengers: the number, taken alone, describes nothing useful about the airport required to serve them. Ten million passengers distributed through narrow-body short-haul operations create one airport: a machine characterised by high gate turnover, rapid apron cycles, dense boarding events, low premium amenity demand and short ground times. Ten million passengers concentrated into wide-body intercontinental banking operations create a fundamentally different airport: one characterised by extended gate occupation, high immigration wave intensity, premium lounge pressure, complex baggage sorting, high fuel demand and long-haul recovery requirements. These are not variations of the same airport; they are different airports that happen to share a passenger total: the design implications — stand geometry, pier width, pushback interaction zones, taxiway occupancy, hydrant loads, catering access, immigration processing, lounge floor area, baggage system throughput — differ in kind, not merely in degree.

Fleet mix compounds the difficulty further: wide-body aircraft do not merely demand more space; they demand different space, deployed differently in time. A 777 at a contact stand requires a pushback that affects adjacent stands; its fuel demand arrives as a volume load that the hydrant system must absorb simultaneously across a bank; its catering requires high-loader access that may conflict with ground-handling equipment for adjacent narrow-body turns; its passengers arrive as a single large surge, creating immigration and security demand that a sequence of narrow-body departures would distribute more evenly. None of these consequences is visible in a passenger total: all of them are visible in an airport designed for the wrong mix.

Annual passenger numbers can therefore remain below forecast while peak-hour failure becomes acute: a circumstance that confounds both parties and requires the tribunal to understand what kind of capacity was actually promised. A building may be empty most of the day and intolerable for two hours; airports live in those two hours, and it is in those two hours that the design either performs or fails.

The deeper question in arbitration is accordingly not whether the forecast was wrong — forecasts are frequently wrong, and the law does not require prophetic accuracy: the question is whether the risk of wrongness was allocated, and if so, to whom, and in what respects. Did the designer merely translate a given forecast into a competent design, leaving the employer to bear the risk that the forecast would prove false? Did the contractor undertake to validate the operating concept, or merely to build to a design it had no authority to question? Did the employer retain control of airline strategy while purporting to transfer design risk? Did the contract define peak-hour conditions, or merely annual volume? Did it distinguish origin-and-destination passengers — who arrive, are processed and depart — from transfer passengers, who arrive already airside and must connect within a fixed window? Did it specify wide-

body stand demand, or merely aggregate stand count? These questions rarely yield to a single clause; their answers emerge from the whole contractual and operational matrix; but they cannot be asked at all until the hidden airline has been identified.

B. Lounges, Hierarchy, and the Passenger the Airport Imagines

The lounge is not an amenity: it is a confession of the passenger the airport imagines, the airline it has designed for, and the commercial theory on which the masterplan depends.

A premium long-haul hub requires the premium lounge not as a luxury but as an operational necessity: space for passengers to absorb connection times of an hour or more without deteriorating into the gate area; showers, sleep rooms and workspaces that convert wait into productive time; security and immigration channels that recognise alliance status; direct flows to gates that allow boarding without the passenger re-entering the retail concourse. Remove these, and the premium passenger's experience collapses — and with it the airline's willingness to designate the airport as a hub for its most profitable traffic: the lounge is not separable from the airport's commercial proposition; it is part of the proposition's physical infrastructure.

The low-cost terminal makes the opposite confession: it imagines price-sensitive passengers, travellers who accepted a departure time they would not otherwise have chosen in exchange for a fare that reflected the savings of constrained amenity. A premium lounge in a low-cost terminal is not merely wasteful; it is a category error — a misunderstanding of the passenger the terminal was built to process. Neither model is intrinsically right or wrong: each is right only if the airport's commercial and operational thesis is right. The design freezes that thesis in concrete and glass: when the thesis proves wrong — when the anchor carrier migrates to a competing hub; when premium traffic fails to materialise at the projected volume; when the hub ambition is quietly abandoned for point-to-point operation — the fixed asset remains as evidence of a judgment that aviation volatility overtook.

Lounges also alter proof in ways that are not immediately obvious: if premium infrastructure is said to have been over-provided, the issue may not be square metres of lounge floor area but a failed forecast of premium traffic, alliance migration, route structure or carrier ambition. If lounges are said to have been under-provided, the employer may argue that the designer failed to appreciate the airport's intended status as a long-haul hub; the designer may answer that the contract never warranted the airline's commercial aspirations: that it was asked to design a terminal, not to guarantee that the national carrier would develop a premium intercontinental network. Both arguments are coherent: the lounge dispute then becomes a dispute about the business model embedded in design — whose assumption was it that premium traffic would materialise; and who bore the risk, in contract, that it would not.

This dynamic extends to every form of passenger hierarchy: fast-track security, premium immigration, business-class routing through the terminal, remote stand avoidance for long-haul connecting passengers, lounge adjacency to long-haul gates, retail sequencing and boarding gate allocation. Every airport contains such a ranking; some announce it; others disguise it as circulation planning. In either case it is a prediction about the passenger population the airport will serve: when that population differs materially from the prediction — because the mix of carriers changes, because the expected hub strategy is abandoned, because low-cost carriers colonise a terminal designed for full-service operations — the hierarchy encoded in the design becomes evidence of a mismatch between the airport built and the airport operated.

III. Procurement: Who Owns the Assumption?

Procurement is commonly described as the allocation of risk; in airport design disputes, it is more accurate — and more illuminating — to describe it as the allocation of assumption: the determination of who is exposed when the predictions embedded in the design prove false.

The distinction is not merely semantic. A risk is usually conceived as a discrete event: a ground condition that exceeds the geotechnical report; a price escalation beyond the contractual allowance; a delay caused by an event outside the contractor's control. The law has developed relatively confident mechanisms for allocating such events. An assumption is more elusive: it may concern how a particular airline will schedule its fleet; how transfer passengers will behave under irregular operations; how security rules will evolve over the airport's operational life; how much dwell time the commercial model can extract from a connection passenger; whether the region will emerge as a tourism destination; whether the national carrier will survive; whether the climate will remain within the tolerances assumed in the pavement design. These are not events that will or will not occur; they are conditions that will or will not persist: the contract rarely identifies them explicitly, and procurement structure determines, imperfectly and often unknowingly, who is exposed when they prove false.

A. Design-Bid-Build and the Refuge of Execution

Design-bid-build appears to offer the clearest allocation: the employer defines; the designer designs; the contractor builds. This works tolerably where the design object is stable: the airport seldom is.

By the time the works reach tender, the employer has typically committed to a constellation of propositions — a political opening date, an anchor airline, a traffic forecast, a concept of operations, a retail philosophy, a security posture, a public narrative — none of which appears explicitly in the bills of quantities. The contractor prices execution, not aviation theory: it undertakes to build what it is shown, not to validate what the design assumes. The assumption, in the design-bid-build model, formally belongs to the employer; but the contractor is the one who assembles the airport's physical parts, integrates its systems, sequences its handover and conducts its ORAT.

When the airport underperforms, the contractor's first line of defence is therefore geometrically attractive: it built the airport it was instructed to build. The employer's response is equally available: execution affects operation; integration, sequencing and handover are not mechanical processes but skilled activities that may themselves have contributed to the airport's failure. Neither answer is complete: a contractor is not ordinarily the guarantor of the employer's aviation ambition; but neither is it merely a mason of documents who owes no duty of skill or care in the manner of its delivery. Airport construction requires coordination across disciplines that no single drawing can prescribe: the point at which execution ceases to be mere execution and begins to affect operation is precisely the point that the design-bid-build allocation of assumption fails to identify.

B. Design-and-Build and the Illusion of Single-Point Certainty

Design-and-build promises simplicity: one counterparty; one contract; one party answerable for the airport. The employer wants a functioning airport and a single throat to hold responsible if it fails: the phrase is crude; the expectation is real; the delivery almost never matches it.

The difficulty is that the employer almost invariably retains the airport's soul while purporting to transfer its body: it prescribes the traffic forecast; the passenger experience standard; the terminal geometry; the airline requirements; the security philosophy; the retail ambition; the expansion strategy; the environmental targets. The design-builder receives these as employer's requirements and is told to develop, detail and deliver them: it has formal design responsibility; it does not have authorship of the airport's aviation thesis; the thesis was fixed before it arrived.

A dispute then develops between formal responsibility and substantive control — and it is a dispute that slogans about single-point responsibility cannot resolve. The employer asserts: you undertook to deliver an airport fit for purpose; the airport is not fit for purpose; you are liable. The contractor responds: you defined the purpose in terms that made performance impossible, or that material changes in the aviation world made unachievable; the purpose was not mine to define, and I priced only what you described. The tribunal cannot adjudicate this exchange without reconstructing the project's intellectual history: what was assumed; who assumed it; who had authority to test it; who had authority to change it; who priced the risk that it would prove false. Those are the questions that dissolve the apparent certainty of design-and-build and expose the allocation of assumption that was always operating beneath the surface.

C. Management Contracting and the Vanishing Integrator

Management contracting, EPCM and multi-package procurement arrangements generate a different and in some respects more troubling problem: integration without an integrator; a system designed to have owners for every component and responsibility for none of the connections.

Airports are interface machines: perhaps the most interface-intensive built assets in existence. The baggage system depends upon architecture, structural capacity, mechanical installation, software, security screening equipment, power supply, airline messaging protocols, ORAT procedures, gate allocation logic and handler staffing. Passenger circulation depends upon retail layout, way-finding, vertical transport, border control configuration, airline schedule design and security rules. Apron operations depend upon stand geometry, fuel hydrant infrastructure, service road layout, ATC procedures, ground handler contracts, rescue access corridors and weather management procedures. In a multi-package procurement, each of these systems may have been designed, supplied and installed by a different party; each party may have performed its own scope with professional competence; the airport may nonetheless fail, because the failure is in the connections — the interfaces — and the connections belong to nobody.

Airport arbitrations arising from multi-package procurement consequently tend to become disputes about governance: not merely about which party was late or defective, but about who owned the obligation to ensure that the system worked as a system. The answer may be the employer; the project manager; the design lead; the systems integrator; the airport operator; or, most troublingly, nobody in any legally meaningful sense. Where aviation volatility collides with a contractual structure that has distributed but not allocated systemic responsibility, the result may be an airport whose failure is attributable to everyone in general and nobody in particular: a forensic labyrinth from which only the most disciplined analysis can extract a legally operative conclusion.

IV. The Airport Beyond the Terminal

The terminal seduces lawyers because it is visible: it can be photographed, walked through, measured and complained about in terms a tribunal can readily understand. This is its seduction; it is also its danger: the airport's performance is frequently determined by systems that never appear in a publicity image — ATC sequencing, stand shadowing, hydrant routing, fire-service response geometry, pavement behaviour under thermal stress, service-road congestion patterns, low-visibility operating procedures, de-icing logic, drainage behaviour under peak precipitation, power redundancy design, security zoning constraints. The real airport is often hidden in its constraints: and the constraints are frequently external to the terminal, beyond the boundary fence, and sometimes beyond the boundary of the contract.

A. ATC, Slots, and the External Limits of Capacity

Runway capacity is not a property of the runway: it is a property of the airspace.

A runway is a strip of pavement. What determines how many aircraft can use it per hour is not its length or surface quality but the separation standards applied by ATC, the mix of aircraft types in the traffic stream, the wake turbulence categories in use, the weather conditions that reduce visual separation to instrument procedures, the staffing levels and equipment of the ATC unit, and the sequencing discipline

maintained across the approach and departure paths. A runway may support, in theory, sixty or more movements per hour; the surrounding airspace, with its departure routes, approach fixes, sector boundaries and noise abatement procedures, may deliver forty-five: that gap is not a defect in the runway, not a defect in the terminal, but it directly affects stand occupancy, gate availability, transfer reliability and baggage processing, because aircraft that arrive late from a compressed sequence occupy gates longer than the design assumes.

Airport design disputes may therefore systematically mislocate causation: gate congestion may appear to reflect inadequate stand planning when the actual cause is ATC-induced compression of the arrival schedule. Conversely, ATC constraint may be invoked too readily as an external excuse for a stand plan that was always insufficient. The tribunal must resist both temptations: the temptation to accept the most visible explanation, and the temptation to accept the most convenient one. Capacity is not a place: it is an interaction between the designed airport and the aviation environment in which it operates.

Slot coordination introduces a further dimension of volatility. At coordinated airports, the slot is the fundamental unit of airport access: a permission to use the runway, the stand and the terminal infrastructure at a specified time. Slots are finite; they are politically and commercially significant; they are subject to use-it-or-lose-it rules that interact directly with operational performance. A design that generates systematic delay will cause airlines to miss slots; missed slots alter stand occupancy patterns; altered stand occupancy disrupts gate assignment; disrupted gate assignment cascades into baggage, connection times and crew scheduling. The airport designed to process sixty movements per peak hour does not merely fail to process sixty movements when constrained: it fails differently in each downstream system, according to which system first absorbs the delay and how quickly it can recover. Slot coordination turns the airport from a physical site into a time market: and time markets, as every derivative trader knows, fail in ways that physical markets do not anticipate.

B. Fuel, Fire, Rescue, and the Hidden Choreography of the Apron

Fuel is one of the most underrated design variables in airport disputes: it appears to be an operational matter and is often, in truth, an architectural one.

The hydrant system — its layout, its capacity, its redundancy, its relationship to stand geometry — determines how quickly aircraft can be fuelled and therefore how quickly they can depart. A hub operation dependent upon tight bank windows cannot treat fuelling as a detail to be resolved by the operator after handover: it is a schedule constraint as fundamental as gate availability, and its failure generates delays as surely as a closed gate. The same logic applies to catering: for a short-haul, low-cost operation, catering is minimal and its impact on turnaround correspondingly small; for a long-haul hub, catering is industrial logistics — bonded stores, cold-chain requirements, high-loader movements, waste removal, security screening,

simultaneous access to multiple wide-body aircraft during a bank — and its integration into the apron choreography is a design responsibility that cannot be discharged by drawing the catering truck routes as an afterthought. Turnaround reliability is the product of how well all these systems — fuel, catering, ground handling, pushback, stand allocation — have been coordinated in the design: when they have not, the failure looks operational because it manifests in operations; its origin, however, is often in design.

Rescue and fire services deserve particular attention because they are the one apron function that the designer instinctively treats as a compliance matter rather than an operational one: vehicle numbers, response times, access routes, water supply, foam provision are prescribed by regulation; meet the regulation; satisfy the requirement; move on. That approach, taken alone, may produce an airport that satisfies its certification criteria and still generates operational fragility: rescue routes that conflict with ground service vehicle movements during normal operations; foam tender positions that restrict pushback geometry; emergency access roads that create constraints on apron layout precisely at the points of highest traffic density. Compliance with regulatory minima is not the same as integration with operational reality: it is in the gap between those two standards that airport design disputes frequently lodge.

C. Pavement, Climate, and the Legal Life of Materials

The runway is a legal problem disguised as a pavement: the choice between asphalt and concrete, or between different asphalt specifications, is presented in design documents as a technical judgment; it is in reality a life-cycle judgment that encodes assumptions about climate, traffic, maintenance windows, closure tolerance and the airport's financial model — assumptions that may persist, or fail to persist, across decades of operation.

Asphalt is generally repairable in shorter time windows: a damaged section can be removed and replaced overnight in many cases, allowing a return to service without extended runway closure. Concrete offers longer theoretical design life and different structural properties under sustained heavy loading. The choice between them depends not merely upon engineering preference but upon the airport's maintenance philosophy, its tolerance for closure, the availability of maintenance windows given traffic patterns, the anticipated aircraft weight and frequency, the climate at the airport site, and the long-term operating model: an airport that cannot close its main runway for extended maintenance without diverting its network carrier must weigh pavement choice differently from a regional airport with a single daily service that can be rescheduled. These are not engineering questions dressed in legal clothes: they are legal questions — questions about risk allocation, about what the designer was asked to optimise for, and about who bears the cost when the optimisation proves wrong.

Climate renders all of this prospective in the deepest sense: airports are long-lived, and runways designed today will be in service in 2060 and beyond. The climate

parameters that informed the design — peak pavement temperatures, freeze-thaw cycle frequency, drainage requirements under extreme precipitation, thermal gradient stresses — may not describe the conditions the pavement will actually experience across its life. One view holds that professional competence in long-life infrastructure design requires engagement with foreseeable climate trajectory, not merely historical climate data: that a designer who ignores the direction of travel cannot claim the safety of the status quo. Another view holds that liability cannot be constructed retrospectively from projections that were uncertain or contested at the time of design: that holding a designer to a climate standard that was not the professional norm at procurement is precisely the hindsight that the law should discipline. The article takes no normative position on that debate; its analytical contribution is narrower: pavement disputes in airport cases illustrate, with unusual clarity, the general problem — the contract seeks to fix a set of physical standards at a moment in time; the world to which those standards must perform extends far beyond that moment; and the gap between the fixed standard and the evolved world is the terrain of the dispute.

V. Dwell, Minutes, and the Opening That Is Not an Opening

A. The Commercial Airport and the Paradox of Dwell

The modern airport earns its profit when the passenger stops moving: a proposition that would have seemed eccentric to an earlier generation of airport planners and should still give pause to those who encounter it for the first time.

The airline's interest and the airport operator's interest are not aligned: the airline wants velocity — rapid turnaround, swift passenger processing, minimal connection time, fast boarding; every minute a passenger spends stationary in the terminal is a minute the airline cannot use to schedule a further rotation. The airport operator, by contrast, depends increasingly upon non-aeronautical revenue: retail, food and beverage, advertising, parking, hotel, car hire, duty-free; non-aeronautical revenue is generated by stationary passengers, and the longer a passenger dwells, the more the airport earns from sources that its aeronautical charges may not capture. The interests are structurally opposed: and the terminal corridor — its width, its retail configuration, its way-finding logic, its gate adjacency — is the physical expression of which interest has prevailed, or of the compromise that neither fully serves.

This opposition is not merely an economic curiosity: it is a design constraint. A terminal corridor too narrow for its retail ambition creates congestion that impairs operational performance; a corridor too wide, or too empty of commercial stimulus, fails the operator's revenue model. Retail positioned at the transfer security checkpoint forces the connecting passenger to navigate commerce when the schedule requires speed; retail positioned after the security checkpoint captures dwell but may be bypassed by the passenger in a hurry. There is no neutral design: every spatial decision takes a position in the conflict between velocity and dwell, and the position taken encodes the airport's commercial thesis as surely as its aviation thesis.

In arbitration, this paradox generates a category of claim unfamiliar to practitioners more accustomed to delay and defect: the claim that the airport failed not to move passengers but to monetise them. Retail frontage may be poorly located; passenger dwell may concentrate in the wrong zones; transfer passengers may bypass commercial areas entirely; congestion may suppress spending by creating anxiety about connection times rather than generating retail browsing. Conversely, a commercially aggressive layout may have impeded operational efficiency sufficiently to generate the very congestion that deterred spending: the dispute then concerns not defect in the engineering sense but failure of the commercial proposition embedded in the design — a proposition that, depending upon the procurement structure and the employer's requirements, may or may not have been the designer's responsibility to validate.

B. Minimum Connection Time and the Tyranny of Minutes

Minimum connection time is one of the purest expressions of airport design as timing: a number that appears to belong to the airline's scheduling department and in fact belongs to the building.

The minimum connection time — the shortest interval, published by the airline and accepted by the slot coordinator, within which a passenger can transfer from an arriving flight to a departing one — is written into the airport's physical geometry. It is a function of pier length: how far does the arriving passenger walk to reach the transfer security point, and from there to the departing gate? It is a function of vertical circulation: how many lifts, escalators and stairs must be used, at what capacity, at what time of day? It is a function of transfer security lane provision and throughput rates; of way-finding legibility; of stand allocation policy; of gate call discipline; and of the baggage system's ability to route a connecting bag to the correct make-up position within the same window. Minimum connection time is the temporal fingerprint of all these decisions combined: it is not a scheduling variable that exists independently of the building; it is a property of the building, expressed in minutes.

The consequences of getting it wrong are not proportional to the error: a ten-minute increase in minimum connection time may appear modest in the abstract; in a banked hub operation, it may determine whether a connecting pair of routes is commercially viable; whether an airline will continue to sell connections over the airport at all; whether the hub generates the transfer volume that was assumed in the traffic forecast on which the design was based. A ten-minute miscalculation may, through this chain of consequences, undermine the airport's entire commercial rationale: the minute is not small merely because it is numerically small; in an airport, minutes have architecture, and architecture has consequences that compound.

Baggage is the airport's nervous system and the most reliable revealer of what minimum connection time actually means in practice: a passenger may make a connection that the baggage system cannot. The passenger walks, clears transfer security, reaches the gate and boards; the bag, meanwhile, navigates a physical and

digital routing system that may have been designed for origin traffic and not for transfer, or for steady flow and not for the surge of a bank. The bag is late; the flight is held or it is not; the passenger arrives without luggage; the airport's transfer performance statistics record a success that the passenger's experience contradicts. Proving a baggage claim requires more than jam statistics or a count of late bags: it requires reconstruction of the chain — where the bag entered the system; what message accompanied it; which routing it followed; where it waited; what make-up position received it; when the aircraft closed. The respondent's answer cannot be merely that each component functioned: the issue may be the choreography, not the instruments.

C. Completion, Certification, and the Opening That Is Not an Opening

An airport may be complete without being ready: a distinction the law must recognise even when the parties prefer to ignore it.

Construction completion, systems completion, regulatory certification, operational readiness, airline migration and commercial opening are different states, achieved at different times, by different processes, under different contractual and regulatory regimes: they are frequently collapsed in argument into the single question of whether the airport was ready by the relevant date; they should not be. A terminal may be physically complete — all structural work done, all finishes applied, all mechanical and electrical systems installed — and yet be operationally immature: its systems not yet integrated, its ORAT not yet complete, its staff not yet trained to the level that actual operation requires. A baggage system may be mechanically complete — all belts turning, all sorters functioning, all screening machines passing their factory acceptance tests — and yet not stabilised to the point where it can reliably process the mix of origin, transfer and transit bags that a real operating day will present.

Opening dates intensify the difficulty because airports are public events with political and commercial significance that the law is ill-equipped to ignore but poorly positioned to accommodate: governments announce opening dates years in advance; airlines publish schedules; ground handlers hire staff; retailers fit out shops; the public is invited to attend a ceremony. By the time the date arrives, the momentum is often irreversible — not because the airport is ready, but because the alternative — a delay — is politically, commercially and reputationally more expensive than proceeding with an airport that is not: the airport is driven toward opening by forces that have nothing to do with contractual readiness, and the first weeks of operation then expose the gap between the opening that was politically necessary and the readiness that the contract required.

In arbitration, the announced opening date is a piece of evidence to be handled carefully: it may demonstrate the employer's expectation; it may demonstrate political pressure; it may demonstrate reliance by third parties; it may demonstrate nothing about contractual completion. The tribunal must ask a more disciplined question: what, precisely, did the relevant contractual milestone require? Taking-over?

Operational readiness? Regulatory approval for the intended use? Beneficial occupation in a defined sense? Completion of ORAT to a specified standard? Migration of airlines from an existing terminal? These are not synonyms; they are different states, with different legal consequences; and the tendency to treat them as interchangeable is one of the most reliable sources of avoidable confusion in airport disputes.

Regulatory certification adds a further dimension: approval may be delayed because the design is defective; because documentation is incomplete; because the employer changed its operating procedures after certification was applied for; because the regulatory authority's safety philosophy evolved between design and handover; because third-party systems — ATC, ground handlers, border control — were not ready in time. Each of these causes has different contractual implications: to say only that the regulator delayed the opening is rarely sufficient analysis. Phased opening raises analogous issues: temporary operational procedures may become semi-permanent workarounds; manual baggage processing may mask system instability; airline route adjustments may be characterised as reasonable mitigation by one party and as evidence of design failure by the other. The tribunal must distinguish operational improvisation from contractual performance, and contractual performance from operational readiness: three different things, frequently conflated, each with its own legal consequences.

VI. Proof: Systems Failure, Models, and the Paper Airport

Airport arbitration is the forensic reconstruction of disappointed assumptions: a description that is not a paraphrase of ordinary construction arbitration, but an identification of a distinctive evidential challenge. In an ordinary construction dispute, the proof problem is usually linear: what was required; what was delivered; what is the gap; what did the gap cost. In an airport dispute, the proof problem is systemic: what was the operational proposition; through which system or systems did it fail; how did the failure propagate; and which part of the propagation is attributable to the respondent's breach rather than to the employer's own assumptions, the airline's behaviour, or the expected and excusable turbulence of opening a major new facility. The difference between a linear and a systemic proof problem is not one of degree: it is one of kind.

A. The Failure of Linear Causation

Airports fail gradually; then suddenly; then all at once: a few minutes lost at security; a few more at the baggage make-up; an aircraft held on stand for a gate that has not cleared; a fuel delay that pushes pushback past the slot; a missed slot that destabilises the departure bank; a transfer bank that collapses because the inbound aircraft are late and the outbound flights will not wait. At no individual point in this sequence does any single system fail in a way that would, in isolation, constitute a serious operational breakdown: taken together, the sequence may defeat the airport's business model for

the day; and if it recurs every day at the same peak, it may defeat the airport's commercial proposition permanently.

This systemic character of airport failure creates a proof dilemma with which both parties must grapple: the claimant cannot prove its case merely by showing that the airport did not work; that is the conclusion, not the analysis. It must identify the chain; locate the break or breaks; demonstrate that the break was attributable to design or construction failure rather than operational, airline or external causes; and show that the failure produced legally recoverable loss rather than mere commercial disappointment. The respondent cannot escape merely by showing that each of its subsystems, tested individually, met its specification: subsystem compliance is not system adequacy, and an airport in which every component performs to standard while the whole fails is not a successful airport; it is an airport in which the failure of integration — itself a design responsibility — went undetected until ORAT or operation exposed it.

The map that the tribunal needs is chronological as well as spatial: it is not enough to trace where a passenger walked or where a bag travelled; the decisive question is when. The same pier may be adequate at noon and impossible at the peak of a bank arrival; the same baggage route may process origin traffic without difficulty and fail for transfer traffic, not because it is defective but because the time window is different and the system was not designed with sufficient margin for the transfer case. Airport proof must attach time to place; and it must do so with sufficient precision to allow the tribunal to distinguish the failure that was foreseeable from the failure that was adventitious. Diagrams, if used well, are not illustrations: they are arguments; a stand plan, an apron sequence diagram, a baggage routing schematic or a security lane capacity model may reveal the causal logic of the failure more clearly than any number of paragraphs — but only if deployed in service of a legal proposition, not in service of a general description of the airport's complexity.

B. Modelling Evidence and the Tyranny of Assumptions

Models have authority because they appear impersonal: the output of software; the product of calculation; the judgment of mathematics rather than of advocates. They are not impersonal: they are collections of assumptions that have been arithmetised.

A passenger-flow simulation embeds judgments about walking speed: typically a single figure, sometimes a range, rarely a distribution that reflects the full diversity of the passenger population — families with children, elderly travellers, passengers unfamiliar with the terminal, passengers in a hurry, passengers who stop without warning to check their phones, passengers who read the departure board and change direction at a corridor junction. It embeds judgments about security throughput: rates derived from trials conducted under controlled conditions, with trained screeners and cooperative passengers, at a time of day and year that may not represent the peak demand the simulation purports to model. It embeds judgments about gate call discipline: the assumption that passengers will respond to boarding calls in an orderly

sequence rather than surging at the gate when the deadline approaches. Each of these judgments, compounded across millions of simulated passenger-journeys, produces a result that looks like a forecast and behaves like a hypothesis.

In arbitration, the critical question about a model is not whether it was produced by reputable consultants using certified software: it is what legal role the model played in the project's development. Was it a planning tool, used internally to test design options before a decision was taken? Was it a reliance document, provided to the contractor as a representation of what the airport was expected to achieve? Was it a contractual performance criterion, against which the delivered airport's performance would be measured? Was it a validation of operational feasibility, given to the employer as assurance that the concept of operations was achievable? Or was it an internal aid, produced at an early stage, overtaken by subsequent design development, and never intended to survive as a definitive statement of expectation? The model's content is one question; its role is another; and only the role determines its legal significance.

Actual divergence between model and operation does not automatically prove breach: airports are affected, especially in their early years, by airline behaviour that differs from the modelled assumptions; by staffing levels that have not yet reached design establishment; by passenger populations that are learning the terminal; by security rules that evolved between the modelling date and the operational date; by seasonal demand patterns that compress the opening period in ways the annual model did not anticipate. Early divergence may be immaturity; persistent divergence across a full operational year may be evidence that the assumptions were wrong: the distinction requires care; the tribunal must resist the temptation to treat every divergence between model and reality as proof of design failure, and must equally resist the temptation to treat every reference to operational immaturity as a complete answer to persistent systemic underperformance.

C. ORAT: When the Paper Airport Meets the Real Airport

Operational Readiness and Airport Transfer exercises are the moment at which the paper airport — the designed, constructed, commissioned airport that exists in drawings, specifications, test certificates and programme milestones — first encounters the real airport: the physical environment in which actual airline operations, actual passenger behaviour, actual baggage routing and actual staff performance interact with the built asset under realistic conditions.

ORAT matters because the paper airport is always more coherent than the real one: on paper, passengers walk at the design speed; bags take the optimal route; security lanes process at the rated throughput; staff respond to system outputs as trained; emergency procedures are executed as written. In reality, none of this is guaranteed: passengers walk at their own speed, which varies by demography, by urgency, by familiarity with the terminal, and by the quality of way-finding; bags encounter exceptions that the system's programming did not anticipate; security lane throughput degrades under peak pressure; staff discover that the procedures that worked in training do not work

in the configuration that actual operations create. ORAT is the controlled collision between paper and reality: its value is precisely that it generates this collision before opening, when remediation is still possible, rather than after opening, when it is merely evidence.

The evidential significance of ORAT in arbitration is correspondingly high: it is contemporaneous; it is operational rather than theoretical; it tests systems in combination rather than in isolation; and it may expose incompatibilities between subsystems that were individually compliant but collectively incoherent. An ORAT that reveals that the baggage system cannot process transfer traffic at the required speed, or that the security lanes become a bottleneck at the peak of a bank, or that the way-finding is insufficient to route unfamiliar passengers to remote gates within the minimum connection time, is powerful evidence that the design's operational proposition was not achievable: it is also evidence that must be handled with care. ORAT exercises are conducted under conditions that are, by definition, not yet normal: staff are not fully trained; systems are not fully stabilised; airline operating patterns are not yet established. Early ORAT instability may reflect transient immaturity rather than persistent design failure; the legal question is what kind of instability was revealed, who owned the underlying cause, and whether the instability was of a kind that further training and familiarisation could be expected to cure.

VII. The Expandable Airport: Designing for a Future That Must Arrive on Time

The expandable airport is not a larger airport held in reserve: it is a promise made in concrete about a future that aviation may or may not deliver, on a timetable that nobody controls.

That distinction matters because expansion clauses in airport design contracts are at times framed as technical provisions — future stand counts, reserved pier zones, structural load allowances, land set-asides, utility duct capacities — when they are in reality commercial and operational commitments of the highest order: the contract that requires an expandable design does not merely require a bigger building someday; it requires that the airport built today remain coherent, functional and commercially viable through the transition to the airport that tomorrow demands. It is that continuity — the obligation not merely to accommodate expansion but to survive it operationally — that generates the distinctive technical and legal complexity.

A. The Capacity Figure and the Illusion of Precision

The expansion mandate most commonly encountered in airport design contracts is expressed as a passenger number: design for twenty million; provide for expansion to thirty-five million; ensure the structural envelope can accommodate fifty million annual passengers. That number appears, on its face, to be a precise instruction. It is not: as argued earlier in this article in relation to capacity generally, a passenger total describes nothing useful about the airport required to serve it, and that observation applies with even greater force when the passenger total in question is a future target rather than a present condition.

Twenty million passengers processed through a low-cost, point-to-point network create one airport: one that demands high gate turnover, dense apron cycling, minimal transfer infrastructure and constrained amenity. Twenty million passengers anchored in a long-haul hub operation create a fundamentally different airport: one that demands banked stand availability, premium lounge capacity, through-baggage infrastructure, transfer security and the redundancy that irregular operations in intercontinental banking require. The design obligation to accommodate twenty million passengers, stated without more, tells the designer almost nothing that is operationally useful: it tells it a number, not an airport.

The difficulty is compounded when the capacity figure is a future target rather than a present one, because the aviation world in which that target will be reached is, by definition, unknown. A contract executed today may specify expansion to thirty-five million annual passengers to be achieved in fifteen years: in those fifteen years, the fleet mix may shift; the anchor carrier may change its strategy or its alliance; low-cost carriers may colonise terminals designed for full-service operations; the balance between origin-and-destination traffic and transfer traffic may reverse; security requirements may transform the relationship between terminal area and processing capacity; environmental constraints may impose operational limits that no masterplan assumed. The number thirty-five million will arrive, if it arrives, in an aviation world that the expansion design cannot yet see and that the original design certainly did not anticipate.

This is not merely a philosophical observation about uncertainty: it is a legally operative point about the scope of the expansion obligation. When the design mandate is expressed as a capacity figure, the first question in any expansion dispute is what that figure was understood to mean at the time of contracting — not in the abstract, but in the specific context of the airline environment, the traffic forecast, the operating concept and the commercial model that the masterplan embodied. A capacity figure divorced from those contextual assumptions is not a design brief: it is a number in search of an airport.

B. Granular Mandates and the Problem of Specified Futures

Some expansion contracts avoid the ambiguity of the bare capacity figure by specifying the expansion in greater operational detail: a defined number of contact stands of specified dimensions; a baggage system capable of processing a defined throughput of transfer bags; a security lane count sufficient for a defined peak-hour passenger flow; pier lengths that achieve a defined minimum connection time. These granular mandates appear to solve the problem of the capacity figure by translating it into physical and operational terms. They solve it only partially; they introduce new problems of their own.

A granular expansion mandate fixes not merely the scale of the future airport but its operational character: in doing so, it encodes, in binding contractual form, predictions

about the future aviation world that are at least as vulnerable to falsification as the bare capacity figure, and considerably harder to revise. A mandate that specifies contact stand dimensions for a defined aircraft type assumes that the aircraft type will dominate the expanded operation; if the fleet mix evolves toward larger or smaller aircraft, the stands may be the wrong size for the expansion they were designed to receive. A mandate that specifies baggage throughput for a defined transfer ratio assumes that the balance between origin and transfer traffic will persist; if the hub strategy is abandoned and the airport becomes predominantly origin-and-destination, the baggage infrastructure may be over-specified in the transfer dimension and under-specified in the origin dimension simultaneously. A mandate that specifies minimum connection time assumes that the pier geometry, the security configuration and the stand allocation logic will remain aligned across the expansion; if any one of these changes — because the expansion design differs from the original design's assumptions, because the airline's scheduling discipline changes, because the regulatory environment for transfer security evolves — the minimum connection time the granular mandate specifies may become unachievable regardless of whether the physical infrastructure is compliant.

The granular mandate, in other words, does not eliminate the volatility-permanence tension: it relocates it. Instead of a single, ambiguous capacity figure, the dispute now concerns the intersection of multiple specific predictions, each of which may have failed independently and each of which may have a different contractual owner.

C. The Expanding Airport and the Presumptive Airline Problem, Revisited

The presumptive airline argument advanced in Section II of this article applies with particular force to expansion mandates, and the application deserves to be made explicit rather than left to inference.

Every expansion mandate encodes a prediction about the airline that will operate the expanded airport: the fleet it will fly, the routes it will serve, the transfer ratio it will generate, the premium passengers it will carry, the dwell time those passengers will tolerate, the connection windows they will require. A bare capacity figure encodes this prediction implicitly, in the assumptions that informed the masterplan from which the figure was derived. A granular mandate encodes it explicitly, in the stand dimensions, baggage throughputs, lounge areas and connection times that the detailed specification prescribes. In either case, the expanded airport has that presumptive airline; and in either case, when that airline proves to be a different airline from the one the expansion assumed, the volatility-permanence tension that the original design already embodied is intensified, because the mismatch between the designed airport and the operated airport now extends across two design exercises rather than one.

This intensification has direct consequences for proof. In an ordinary airport design dispute, the tribunal must reconstruct the airline prediction embedded in the original design and determine whether its failure was a design defect or a volatility event

outside the designer's responsibility. In an expansion dispute, the tribunal must perform that reconstruction twice — once for the original design and once for the expansion design — and must then determine whether the failures at the two stages are independent or compounding. If the original design embedded an airline prediction that the expansion design then relied upon, and if that prediction proved false, the question of who bears responsibility for the compounded failure is not resolved by the simple observation that both stages shared the same assumption: it requires analysis of which stage had the authority and the incentive to test, challenge or revise that assumption, and which stage merely inherited it.

D. Structural Expandability, Systemic Expandability, and the Gap Between Them

An expansion mandate typically operates on at least three distinct levels, and the failure to distinguish them generates disputes in which the parties speak entirely past each other.

The first level is structural: the building must be capable of receiving additional floor area, additional piers, additional stands, without requiring demolition of what was built first; load paths must be designed for future vertical extension; foundations must carry loads that the initial building will never impose; structural bays must align with future column grids; roof levels must allow future connection at the right datum. These are engineering obligations, and their breach is usually detectable even when late: a foundation that cannot carry a future column is a fact, and it is discoverable.

The second level is systemic: the airport's infrastructure must be capable of serving the expanded facility without wholesale replacement; baggage system capacity must have headroom; hydrant mains must be sized for future stands; electrical substations must have reserve capacity; ATC infrastructure must accommodate additional movement rates; security lane counts must scale; immigration processing must expand without relocating the entire border control operation. These obligations are harder to assess at handover, because they are latent: they will not fail until the expansion occurs and the infrastructure is called upon to serve it. That latency creates a proof problem of its own — the question of when the breach occurred, and therefore when time began to run, may be genuinely uncertain, because the deficiency that constitutes the breach was invisible until the expansion imposed the load that exposed it.

The third level is operational: the airport must be capable of remaining in operation — fully, or at an acceptable degraded standard — while expansion occurs around it. This level receives the least attention in design contracts and generates the most disputes: phased construction in a live airport is not an engineering inconvenience; it is a sustained operational challenge in which construction activity competes directly with the aviation movements that the airport exists to serve. Scaffolding near taxiways; temporary barriers on aprons; construction vehicle routes across service roads; noise and vibration near ATC facilities; dust on runways; interim baggage routing while

permanent infrastructure is extended — each of these is a point of interaction between the fixed world of construction programming and the unforgiving world of aviation operations, and the contract that fails to address these interactions in operational terms will discover, during the expansion, that the gap between structural expandability and operational continuity is precisely where the most consequential and the most contested failures occur.

E. The Masterplan as Moving Target

The expandable airport is almost always accompanied by a masterplan: a document that attempts to describe not merely the airport as it will be built but the airport as it will become, over a period that may extend to twenty or thirty years. The masterplan is the instrument through which the expansion promise is given architectural and planning form; it is also, in any dynamic aviation environment, a document under pressure from the moment it is published.

Airlines do not grow according to masterplans: route networks expand and contract in response to competition, fuel prices, regulatory change, geopolitical events and the strategic decisions of airline management whose tenure may be shorter than the masterplan's horizon. A masterplan predicated upon the growth of a hub carrier may be overtaken within five years by that carrier's restructuring, alliance change or bankruptcy. A masterplan that assumed steady narrow-body growth may be confronted instead by the rapid introduction of ultra-long-haul twin-engine operations that require different stand dimensions, different fuel loads, different pier widths and different immigration processing capacity than the plan assumed. A masterplan that anticipated a second runway within fifteen years may find that environmental constraints, community opposition or regulatory delay push that timeline to thirty, leaving the single-runway airport to absorb traffic volumes that were supposed to be distributed across two.

Each of these scenarios transforms the expansion obligation without changing its text: the contract still says thirty-five million passengers; the masterplan still shows a second pier; the expansion clause still prescribes the stand count. But the airport that those provisions were designed to produce no longer corresponds to the airport that aviation reality will require. Whether the divergence between the masterplan's expansion trajectory and the actual trajectory of aviation growth constitutes an employer's variation, a change in the basis of design, a risk that the designer always carried, or an event outside any party's contractual responsibility, is a question the contract will rarely answer directly: it is a question that the tribunal must reconstruct from the procurement structure, the applicable law, and the degree to which the original designer had authority or incentive to challenge the masterplan assumptions at the time of contracting.

F. Advocacy and Determination in Expansion Cases

The distinctive complexity of expansion disputes has direct consequences for both advocacy and determination.

For the advocate, the first discipline is temporal: expansion disputes involve at least two distinct periods — the original design and construction, and the expansion itself — and potentially a third, the operational transition between them; claims and defences that conflate these periods will almost always be wrong, because the obligations, the parties and the applicable assumptions differ across them. A rigorous expansion case requires a timeline that is not merely chronological but jurisprudential: identifying, at each stage, what obligation applied, who owed it, what the relevant assumption was at the time, and whether the assumption was the same assumption that the preceding stage had made or a new assumption that the expansion design introduced.

The second discipline is the capacity figure: wherever the expansion mandate is expressed as a passenger total, the advocate must translate that total into the operational terms that give it meaning. What type of passengers? What fleet mix? What transfer ratio? What peak-hour distribution? What connection time? What commercial model? The answers to these questions determine what the expansion obligation actually required; and the answers are almost never found in the expansion clause itself; they must be reconstructed from the masterplan, the employer's requirements, the traffic forecast, the operating concept and the commercial assumptions that informed the design at the relevant stage. An advocate who accepts the capacity figure at face value and argues about whether the expanded airport meets it, without first establishing what the figure was understood to mean, has conceded the most important issue in the case before the hearing begins.

For the tribunal, the determination of an expansion case requires a framework capable of holding the temporal and systemic layers of the dispute simultaneously without collapsing them into each other: the instinct to simplify — to find a single cause, a single obligation, a single period — will be wrong in most expansion disputes, because the disputes are structurally multilayered and the award must be too. That does not mean the award must be long; it means it must be precise about which layer of the dispute it is addressing at each point in the analysis, and about how the conclusions at each layer interact with those at others.

The expandable airport is the volatility-permanence tension raised to its highest power: an attempt not merely to fix the present in concrete, but to provide in the present for a future that aviation will define on its own terms and in its own time. The capacity figure that sits at the heart of the expansion mandate — twenty million, thirty-five million, fifty million — is not a specification; it is a hypothesis about an aviation world that does not yet exist: and it is the tribunal's task, when that hypothesis fails,

to determine with discipline and precision who, in contract as understood by reference to the applicable law, agreed to bear that risk.

VIII. Law: Standard of Care, Applicable Law, and the Frontier of Loss

The law likes classifications: means or results; contract or tort; direct or consequential loss; foreseeable or remote. Airports unsettle them all — not because they are legally mystical, but because they combine, in a single physical and contractual object, public purpose and private commerce; technical engineering and human behaviour; professional judgment and operational outcome; sovereign ambition and financial constraint. The volatility-permanence tension that runs through every aspect of airport design reappears in legal form as a single, insistent question: what, exactly, did the designer or contractor undertake to do about a future that could not be fixed?

A. Standard of Care and the Operational Undertaking

The standard of care in an airport case is almost always stated too quickly: the first statement — that the designer must exercise reasonable skill and care — is typically correct as far as it goes; the difficulty is that it does not go far enough to resolve the dispute.

The designer will say: professional competence cannot guarantee the future; airport design is prediction under constraint; the variables — airline strategy, passenger behaviour, security evolution, climate, ATC constraints, political ambition — are beyond the designer's control; and liability indexed to those variables is liability that no professional can price or insure. That argument is coherent: a designer cannot be the insurer of aviation's unpredictability.

The employer will respond: the designer was not retained to draw an attractive terminal abstracted from aviation reality; it was retained precisely because of its expertise in the aviation environment that the terminal was designed to serve; that expertise consists in understanding, anticipating and managing the very volatility that the designer now invokes as a defence; and a design that ignores the instability that defines airports cannot defend itself by pointing to that instability when it materialises. That argument is also coherent: professional skill in airport design is not merely skill in producing compliant drawings; it is skill in producing designs that will work in the aviation world, under the conditions of uncertainty that characterise that world.

The resolution — to the extent that any general resolution is possible — lies in the specific contract, at times as supplemented or even supplanted by the applicable law, and the particular allocation of assumption that the procurement structure created: did the contract prescribe the traffic assumptions, relieving the designer of responsibility for their accuracy; did the designer have authority to challenge the concept of operations, or was it required to develop a concept it had no power to alter; were the traffic forecasts provided as employer information on which the designer was entitled to rely without independent validation; were performance criteria framed as design inputs — the conditions the designer was told to design for — or as operational

outputs — the results the designer was required to achieve? The same professional conduct may constitute breach under one analysis and compliance under another; and the analysis cannot be performed without understanding the airport well enough to identify what the designer knew, what it should have known, and what it had authority to do about what it knew.

Expert evidence on standard of care requires corresponding discipline: the expert must resist judging from hindsight; the question is not whether the design now appears inadequate in the light of what the airport became, but whether a competent airport professional, working at the relevant time with the relevant brief, information and contractual authority, would have produced a materially different design. The expert must equally resist the temptation to excuse thin design by invoking the general uncertainty of aviation: uncertainty is the airport's defining condition; it cannot simultaneously be the environment that creates the need for specialist expertise and the defence against liability for the exercise of that expertise. A good award will ask not whether the airport now disappoints, nor whether the designer's documents were produced in the conventional sequence: it will ask whether the professional obligation, properly characterised in the light of the contract and the applicable law, required the party to identify, test, challenge, warn against or absorb the operational instability that later materialised.

B. Contract, Applicable Law, and Ex-Contractual Routes

International airport projects present the law with a compound challenge: an elaborate contract that is nonetheless incomplete; an applicable law that may not have developed its principles in the airport context; and a set of operational consequences that may exceed what either the contract or the applicable law was designed to address.

The contract may specify capacities, deliverables, testing regimes, performance milestones and limitation of liability provisions with considerable sophistication; it may nonetheless fail to address the assumption that later proves critical: the assumption about how the airline will schedule; about how the employer will operate the terminal in the transition period; about what changes to security rules will be tolerated within the contracted scope; about how the applicable law's implied terms of professional care or fitness for purpose will interact with the express allocation of risk. When the contract is silent or ambiguous on the critical point, the applicable law enters: its rules of interpretation; its implied terms; its doctrines of good faith, cooperation, professional duty, reliance and remoteness; its understanding of gross mistake and deceit; its equitable principles of unjust enrichment; its approach to concurrent causes; its treatment of pure economic loss.

Different legal systems answer these questions differently, and the difference is material in airport cases precisely because the legally critical questions — what was implied; what was foreseeable; what was assumed — are questions to which the aviation context gives highly specific, and often counterintuitive, answers. International airport projects frequently mix drafting traditions: a FIDIC-derived

contract may be administered under a civil-law system with strong public-law overlay; the expert evidence may follow common-law traditions; the arbitration procedure may be international; and the damages analysis may import concepts from a third legal system. The resulting grammar is rarely pure: and the analysis must account for the fact that terms familiar across legal systems may carry different freight in each.

One of the reasons ex-contractual claims arise in airport cases is that the airport's operational consequences routinely exceed the immediate contractual relationship: airlines, passengers, concessionaires, ground handlers, regulators and regional economies may all be affected by airport underperformance; none of them is a party to the design contract; and the question whether the designer owes them duties in tort, delict or professional responsibility — duties that survive alongside or beyond the contractual allocation of risk — is one that airport cases raise with unusual force. The answer depends upon jurisdiction and facts. But the question cannot be avoided in airport disputes because the operational radiation of design decisions extends far beyond the bilateral bargain, and the law must decide how much of that radiation it is willing to recognise.

C. Strategic Economic Loss and the Identity of the Employer

The airport is rarely sold as infrastructure alone: a tourism engine, a logistics platform, a convention gateway, an investment signal, a generator of employment, an instrument of economic diversification, a statement of national ambition. These characterisations appear in political speeches, in procurement documents, in employer's requirements, in project narratives prepared for lenders and in the environmental impact assessments that underpin planning approval: they are not decorative; they describe the purpose for which the airport was procured and, in some cases, the basis on which its financing was structured.

When the airport underperforms, alleged losses may therefore move beyond aeronautical revenue and retail income to encompass lost tourism, reduced hotel occupancy, diminished convention traffic, frustrated inward investment, impaired cargo growth and the broader economic consequences of aviation underperformance in a regional or national economy. These claims are not frivolous: in jurisdictions or contexts where the employer procured the airport as an instrument of economic development, the losses that flow from its failure may be exactly the losses that the employer was created or mandated to avoid. But they raise questions of causation, remoteness, proof and recoverability that strain the ordinary categories of construction damages, and that may produce very different answers depending upon the applicable law and the identity of the contracting parties.

The identity of the employer matters because it shapes the proximity of the claimed loss to the contractual relationship: where the employer is the sovereign state itself, the broad economic consequences of airport underperformance may appear conceptually close, because the state embodies the public economic interest that the airport was designed to advance, and the loss to that interest is at least notionally the

state's loss. Where the employer is a corporatised airport operator, a special-purpose vehicle, a concession company or a municipal authority, the same consequences may appear more attenuated: the loss, if it exists, may belong to the tourism industry, to regional employers, to tax revenues, to third-party investors — to persons who are not party to the contract and whose losses may not be recoverable through it. The identity of the employer thus determines not the existence of the loss but the legal distance between the loss and the party that can sue, especially in arbitration: that distance may or may not be bridgeable; but the question is unavoidable in any airport dispute that involves an employer with a public or developmental mandate, and it resists resolution by the formulae developed for ordinary commercial construction.

IX. From Airport to Award: The Tribunal's Task

Although these cases tend handsomely to repay immersion in the technical detail (if one is so inclined), the tribunal's task is not to become an airport expert but to understand enough of the airport lifecycle to decide the legal question correctly. That formulation sounds modest; it is not: the gap between understanding enough and understanding too little is precisely where airport awards go wrong; and closing that gap requires more than a general familiarity with airports derived from the experience of being a passenger.

Airport disputes tempt tribunals toward two opposite errors, each as damaging as the other. The first is abstraction: treating the case as a variant of familiar construction arbitration and pressing the airport's distinctive features into categories — delay, defect, variation, quantum — that were not designed for them; this produces awards that are legally tidy and operationally illiterate: awards that resolve a building dispute when the parties brought an airport dispute. The second is immersion: disappearing into the operational complexity of the airport until the legal question loses its shape; producing an award that demonstrates the tribunal's mastery of baggage system throughput rates and ATC sequencing procedures without ever clearly stating what obligation was breached, by whom, with what causal consequence and at what legally recoverable cost. The right course lies between these errors: the tribunal must be fluent enough in airport operations to follow the argument, and disciplined enough in legal analysis to resist the seduction of operational complexity as a substitute for legal reasoning.

Pleading discipline follows from the same logic: a pleading that asserts merely that the airport was under-designed, delayed or defective is too blunt to found an airport claim; it states the conclusion without identifying the operational proposition said to have been promised; the contractual route by which that proposition became a binding obligation; the specific subsystem or interface at which the obligation was breached; the operational consequence of that breach in the airport's functioning; and the legally recoverable loss that followed as a matter of causation and applicable law. Omit any one of these steps, and the claim becomes atmospheric: a complaint about the airport's performance dressed in legal language, rather than a legal claim arising from a specific breach of a specific obligation. The respondent's discipline is correspondingly

demanding: it is not enough to deny that the airport failed, and not enough to recite compliance certificates; the respondent must identify the design basis; the assumptions that the employer supplied; the later departures from those assumptions; the intervening operational causes that contributed to any apparent dysfunction; and the principled distinction between technical performance of the contracted scope and commercial disappointment attributable to aviation volatility that the contract did not allocate to the respondent.

Expert casting in airport cases is an exercise in selection: the airport's systemic character may appear to require expert evidence on every system — passenger flow, baggage, ATC, pavement, fuel, rescue, retail economics, ORAT, security, airline operations; in practice, a case that presents all of these simultaneously may illuminate none of them adequately; the volume of expert material will sometimes exceed some tribunals' capacity to absorb it; the experts will not have been given a sufficiently focused mandate to produce evidence that serves the legal issues; and the forensic exercise will degenerate into a competition between extensive expert reports rather than a disciplined reconstruction of the operational failure. The art is to identify, from the airport's systemic complexity, the operational chain that is legally significant — the specific sequence of events, through the specific systems, that connects the alleged breach to the alleged loss — and to instruct experts who serve that chain rather than describing the airport in its entirety.

The structure of the award should reflect the same discipline: where relevant and appropriate, it should identify the airport's operating proposition — what the airport was designed to do, for which airline and passenger population, under which operating concept — and determine whether, and through which contractual route, that proposition became a binding obligation. It ought to identify the relevant assumptions: who supplied them; who was entitled to rely on them; who had authority to test or change them; and who bore the contractual risk of their failure. It should then address breach, causation and loss in that order, resisting the temptation to allow operational evidence to substitute for legal analysis at any stage.

The best airport analysis proceeds backwards: from the operational failure to the assumption it exposed; from the assumption to the party that owned it; from the ownership to the legal obligation; from the legal obligation to the breach; from the breach to the causally connected loss. This method does not predetermine results: it disciplines inquiry. It prevents the claimant from converting every operational disappointment into breach by requiring it to identify the specific contractual promise that was not kept; it prevents the respondent from hiding behind technical compliance by requiring it to confront the possibility that what was promised was not technical performance but operational adequacy — possibly a higher and different standard; and it prevents the tribunal from resolving the case at the wrong level of abstraction: neither so abstract that the airport's operational character is ignored, nor so granular that the legal question is obscured by the details of the baggage system's sortation logic.

The most persuasive airport advocacy will therefore often be the least encyclopaedic: it will not attempt to establish that every operational inconvenience is attributable to the respondent, because that approach disperses forensic attention and invites the respondent to demonstrate that each individual complaint has an innocent explanation. It will identify the decisive time relationships, the decisive assumptions, the decisive contractual promises and the decisive departures from them: it will show the tribunal how the airport was designed to breathe — the sequence of events, from aircraft arrival to aircraft departure, that the operating concept assumed would occur within the required time windows — and then show, with precision, where the rhythm broke; or why, despite appearances, the rhythm's failure was or was not the respondent's responsibility.

X. Conclusion: The Impossible Stability of Airports

The airport seeks to stabilise motion: that is its ambition, and its fragility, and the source of every difficulty examined in this article.

The attempt is necessary: without fixed structures, contracts, operating concepts and time sequences, the aviation system cannot function. The attempt is also never complete: aviation is constitutively unstable, driven by competition among carriers, by the volatility of demand, by the unpredictability of security and political conditions, by climate and by technological change. The contract and the masterplan aspire to fix, for decades, arrangements that aviation reality will continue to revise: they do not always fail because they are poorly drafted or poorly designed; they sometimes fail because the aspiration is in tension with the nature of the thing it seeks to stabilise. Arbitration begins when that tension produces a disagreement about who bore the risk of the gap between the stability that was promised and the instability that prevailed.

The collision between volatility and permanence is not one difficulty among others in airport arbitration: it is the source of all the others, and each of this article's themes is an expression of it. Capacity becomes treacherous because the contract fixed a number while aviation changed the fleet. Procurement fails to settle responsibility because the assumption embedded in the design was never identified clearly enough to be allocated explicitly. Proof becomes a systemic challenge because the failure is the product of an interaction between fixed design and moving operations, and the chain of causation cannot be reconstructed by examining any single link. The standard of care becomes unstable because the professional obligation to manage volatility cannot be assessed without first determining how much volatility the contract required the professional to manage: a question that admits of no general answer. Loss becomes contested because the airport's consequences extend beyond the project boundary into an aviation and economic environment that the contract did not, and could not, fully address.

None of this means that airport disputes resist resolution: it means that they resist resolution by the methods and categories that suffice for simpler construction disputes, and that the tribunal, the advocates and the experts who approach an airport

case with those methods will find, at some point in the proceedings, that the airport does not cooperate. The appropriate response is not to force the airport into the familiar categories: it is to develop the analytical tools that the airport requires.

Those tools begin with reclassification: the recognition that the airport is a synchronised system, not a building; that its failure is typically systemic, not componential; that its obligations are typically operational, not merely technical; and that its losses are typically downstream of the aviation world it was designed to serve, not merely downstream of the construction process that produced it. From that reclassification, the legal analysis follows with greater reliability: the standard of care becomes more precisely statable; the allocation of assumption through the procurement structure becomes more legible; the causal chain becomes more traceable; and the boundaries of recoverable loss become, if not certain, at least more principled.

Airport arbitration begins where operational volatility collides with contractual permanence: that collision is this article's subject. It is also the reason airport disputes cannot be resolved as ordinary construction disputes with aircraft attached: the aircraft, and everything that aviation requires of the ground, changes the nature of the legal problem, not merely its scale. To understand that change is the beginning of the analysis: everything else follows.