



Received: 29-07-2026
Accepted: 09-09-2026

International Journal of Advanced Multidisciplinary Research and Studies

ISSN: 2583-049X

Mathematical Models of Oversupply Risk in the Coastal Upscale Hotel Market of Vietnam: Entry Timing, Construction Lags and an Early Warning Index

¹ Nguyen Thi Thu Ha, ² Nguyen Xuan Bach

¹ Graduate Academy of Social Sciences, Vietnam

² Foreign Trade University, Vietnam

Corresponding Author: Nguyen Thi Thu Ha [Email: bachcu012004@gmail.com]

DOI: <https://doi.org/10.62225/2583049X.2026.6.5.6913>

Abstract

Coastal resort cities in Vietnam have added upscale hotel rooms at a pace that repeatedly outruns demand. The usual explanation, developer optimism, sits awkwardly beside the fact that most of these projects are appraised with conventional discounted cash flow methods and financed by institutions with a strong incentive to avoid loss. This paper develops a mathematical framework in which oversupply emerges as a property of the market rather than as a failure of individual judgement. Three blocks are combined: a delayed stock-and-flow accounting of rooms in preparation, under construction and in operation; a real options trigger governing when the marginal developer starts building; and a composite oversupply risk index that summarises the occupancy gap, the forward absorption time of the pipeline

and the distance between asset value and replacement cost. The balanced growth path is derived analytically. We obtain a closed-form stability condition linking the construction lag to a critical feedback gain, and show that the cycle length at the stability boundary is $2(2\tau + 1)$ years, so that a three-year lag falls inside the ten to fifteen year range reported for property cycles. Simulation under illustrative parameters indicates that competitive erosion of the option premium, rather than forecasting error, is the most powerful single driver of cycle severity, and that a cap on annual construction starts dominates both disclosure and lag reduction. The framework is specified so that it can be calibrated to one coastal city once room supply, pipeline and rate data are assembled.

Keywords: Hotel Supply Cycle, Oversupply Risk, Real Options, Time to Build, Early Warning Index, Destination Planning

JEL Classification: L83, R33, C63, D92

1. Introduction

A coastal city that has just recorded three consecutive years of rising room rates is a city in which hotel construction is about to accelerate. Four years later, when those rooms open, the same city is likely to be reporting falling occupancy, discounted rates and strained refinancing. The sequence is familiar enough that practitioners treat it as a fact of the trade, and it has been documented across property types and institutional settings since at least Wheaton (1987) ^[37]. What remains unsettled is why it keeps happening in markets where the participants can see it coming.

The question has practical weight in Vietnam. Upscale accommodation in the country's coastal resort cities has been built in concentrated waves, each wave following a period of strong rate growth and each delivering into conditions noticeably weaker than those that justified it. Provincial authorities approve projects one at a time, using criteria that are reasonable for any single project considered on its own. Lenders appraise each hotel against its own feasibility study. Nobody in the chain is obviously behaving foolishly, and yet the aggregate outcome is a stock of rooms that periodically exceeds what the destination can fill at rates covering the cost of capital.

This paper argues that the outcome is structural. Three features of the accommodation market interact in a way that produces cycles even when every participant is rational and well informed. Construction takes years, so the supply that arrives today was decided under conditions that no longer hold. Entry is governed by a threshold rather than a smooth response, because a developer with an irreversible investment and an uncertain payoff has value in waiting. And the waiting itself is not idle: while

the threshold is unmet, sites are assembled, designs are drawn and approvals are obtained, so that a stock of shovel-ready projects accumulates and is released in a burst when the threshold is finally crossed. ‘

None of these three features is new. Time to build has been studied since Kydland and Prescott (1982) ^[21] and Majd and Pindyck (1987) ^[24]. The option value of waiting is the subject of a mature literature following McDonald and Siegel (1986) ^[25] and Dixit and Pindyck (1994) ^[11]. The tendency of competition to dissipate that option value was established by Grenadier (1996, 2002) ^[18, 19] and Leahy (1993) ^[22]. What has not been done, to our knowledge, is to put the three together in a model of accommodation supply that is simple enough to calibrate on the data a destination authority actually holds, and then to ask which policy lever does most to damp the resulting cycle.

The contribution is fourfold. We specify a compact dynamic model in which prepared projects, projects under construction and operating rooms are tracked as separate stocks, with an option-based gate controlling the flow between the first two. We derive the balanced growth path in closed form and obtain an analytical stability condition: the linearised system oscillates without damping once the feedback gain exceeds $2 \sin[\pi / (2(2\tau + 1))]$, where τ is the construction lag in years. The same derivation gives a boundary cycle length of $2(2\tau + 1)$ years, which reproduces the decade-and-a-half rhythm reported for property markets without any parameter chosen to that end. We then construct an oversupply risk index from three quantities a planning office can observe or estimate, and finally we rank four policy instruments by their effect on the depth of the cycle.

A word on what this paper is and is not. It is a modelling paper. The parameter values used in the simulations are illustrative, chosen to lie within ranges reported in the literature cited in Section 4, and they are placeholders for an empirical calibration that requires data not yet assembled in a single source for any Vietnamese coastal city. Section 4 sets out precisely which series are needed and where they can be obtained. No result reported below should be read as an estimate of conditions in any actual city, and the paths are demonstrations of model behaviour, not forecasts. We regard this division of labour as the honest one: the analytical results stand on their own, the numerical results show what the model does, and the empirical content waits on the data.

The remainder of the paper proceeds as follows. Section 2 places the argument in the literature. Section 3 sets out the model and the two analytical results. Section 4 describes the calibration strategy and the data an empirical application would require. Section 5 reports the behaviour of the model under deterministic and stochastic demand. Section 6 evaluates policy instruments. Sections 7 and 8 discuss implications and limitations, and Section 9 concludes.

2. Related Literature

2.1 Supply cycles in accommodation and property markets

The stock-flow tradition in property economics treats the market for space and the market for assets as linked but distinct. Rents clear the space market given the standing stock; asset prices capitalise those rents; construction responds to the gap between asset price and replacement cost; and the new stock arrives with a delay. DiPasquale and Wheaton (1992) ^[10] give the canonical statement. Wheaton (1987) ^[37] showed for offices that the delay alone is

sufficient to generate cycles in a model with entirely rational participants, and Wheaton and Rossoff (1998) ^[38] carried the argument into the lodging sector, where they found that the delivery lag and the elasticity of construction to profitability jointly govern the amplitude of the cycle.

Hotels differ from offices in ways that matter for this mechanism. Leases are one night long, so the space market clears almost continuously and revenue responds immediately to a supply shock rather than being cushioned by contracts. That makes the feedback from the stock to the asset value faster than in office markets, which, as we show, pushes the system towards the unstable side of the boundary rather than away from it. Grenadier (1995) ^[17] found that local rather than national conditions dominate vacancy dynamics, an observation that applies with more force to resort destinations, where the demand catchment is defined by the destination itself.

The destination literature has treated the same phenomenon from a different angle. Butler (1980) ^[7] described a life cycle in which a destination passes through exploration, development and consolidation towards stagnation, and Getz (1992) ^[15] linked that trajectory to planning failure. Carrying capacity assessment, reviewed by Saveriades (2000) ^[28] and by Coccossis and Mexa (2004) ^[9], asks how much visitation a destination can absorb before its physical or social base degrades. Both traditions are concerned with too much development, but neither supplies a dynamic mechanism that explains the timing of construction waves, and neither yields an indicator that can be computed annually.

2.2 Investment timing under uncertainty

An irreversible investment with an uncertain payoff should not be made as soon as its net present value turns positive. The right to wait has value, and exercising it destroys that value, so the investment threshold exceeds the replacement cost by a multiple determined by volatility, the discount rate and the payout rate. McDonald and Siegel (1986) ^[25] established the result and Dixit and Pindyck (1994) ^[11] gave the standard treatment. Titman (1985) ^[36] and Capozza and Helsley (1990) ^[8] applied the logic to urban land, and Bar-Ilan and Strange (1996) ^[3] showed that a construction lag interacts with the option value in a way that can lower the trigger rather than raise it, because a project started today delivers into a future the developer cannot wait out.

Applied to hotels, the textbook trigger multiple is large. For plausible volatility and discount rates it lies between roughly 1.8 and 2.8, meaning that an asset should be worth around twice its construction cost before a single developer with a secure claim to the site would rationally break ground. Observed behaviour looks nothing like this. Projects proceed at valuations close to replacement cost, which is what the standard appraisal practice described in the hotel investment literature would recommend.

The resolution, and it is the hinge of our argument, is competition. Grenadier (1996) ^[18] showed that when developers compete for a market that will not support all of them, the fear of pre-emption collapses the option premium and can produce development cascades, with construction starting in a bunch and continuing into a downturn. Leahy (1993) ^[22] established the more general result that competitive entry behaves as though firms were myopic, even though each firm individually holds a valuable option. Grenadier (2002) ^[19] extended the Dixit and Pindyck framework to oligopoly and showed how the premium

erodes as the number of competitors rises. A coastal resort city with a limited stock of beachfront sites and many developers is close to the conditions under which the premium is small. We therefore treat the share of the premium that survives competition as a parameter and trace its consequences.

2.3 Indicators and early warning

Macroprudential research has produced a mature methodology for turning slow-moving imbalances into an index that leads a crisis. Borio and Lowe (2002) [6] and Borio and Drehmann (2009) [5] built on the credit-to-output gap; Kaminsky, Lizondo and Reinhart (1998) [20] formalised the signalling approach; Drehmann and Juselius (2014) [12] evaluated which indicators actually lead. The common design principle is that a useful indicator must be computable from observable data with a lead time long enough to allow a response, and must combine a stock measure with a flow measure so that it responds to commitments already made rather than only to outcomes already realised.

Accommodation markets have no equivalent. Industry practice monitors occupancy, average daily rate and revenue per available room, all of which are contemporaneous, and pipeline counts, which are reported as raw room numbers rather than as a ratio to anything. The index proposed in Section 3.4 is an attempt to close that gap using the macroprudential design principle: one contemporaneous component, one forward-looking component built from the pipeline, and one valuation component.

2.4 Quantitative modelling in tourism and hospitality

Quantitative work in tourism has concentrated heavily on demand forecasting. Song and Li (2008) [33] and Song, Qiu and Park (2019) [34] review a literature that now runs to many hundreds of studies, and the scientometric analysis of Liu, Liu, Wang and Pan (2019) [23] shows that comparisons of forecast accuracy across model families dominate the field. Forecasting demand is necessary but not sufficient for the problem studied here, because the damage in an oversupply episode is done by supply decisions taken years earlier, not by the demand forecast itself.

Operational modelling in hospitality has concentrated on revenue management, surveyed by Bandalouski, Kovalyov, Pesch and Tarim (2018) [2], where the decision horizon is the booking window rather than the building cycle. Efficiency measurement, reviewed by Assaf and Josiassen (2016) [1] and applied to hotels by Barros (2005) [4], asks how well existing capacity is used rather than how much capacity should exist. System dynamics has been applied to tourism planning, as Sedarati, Santos and Pintassilgo (2019) [29] document, but mostly to environmental and social carrying capacity rather than to the accommodation investment cycle. The gap this paper addresses lies between these literatures: a supply-side dynamic model, specified at the level of a single destination, with an entry rule derived from investment theory and an output that a planning authority can act on.

3. The model

The model describes one destination, defined as the market area within which upscale hotels compete for the same visitors, observed annually. Three stocks are tracked. Prepared projects are rooms for which a site has been secured and approvals obtained but construction has not

begun. Rooms under construction are those started and not yet delivered. Operating rooms are the standing stock. Fig 1 shows the structure and the two feedback loops that drive it.




Fig 1: Structure of the model. Solid boxes are stocks and flows that are measured directly; dashed boxes are valuation quantities inferred from them. The balancing loop B1 operates with a delay of τ years, which is the source of the cyclical behaviour analysed in Section 3.5

3.1 Demand, occupancy and rates

Room-night demand grows at a trend rate g , is displaced by a persistent stochastic component and by any scenario shock, and responds to the average daily rate with constant elasticity:

$$D(t) = D(0)(1 + g)^t \exp[u(t)] m(t) [A(t)/A(0)]^{-\eta} \quad (1)$$

$$u(t) = \rho u(t-1) + \varepsilon(t), \varepsilon(t) \sim N(0, \sigma_D^2) \quad (2)$$

Where $m(t)$ captures a scenario shock and equals one in the absence of one. Occupancy follows from the standing stock, subject to a practical ceiling that reflects the impossibility of selling every room on every night:

$$O(t) = \min\{O_{max}, D(t) / [365 K(t)]\} \quad (3)$$

The rate adjusts to the occupancy gap rather than clearing the market instantaneously, which reflects the contractual and reputational frictions that prevent hotels from repricing without limit:

$$\ln A(t+1) - \ln A(t) = \alpha [O(t) - O^*] \quad (4)$$

Revenue per available room is the product of rate and occupancy, and the value of a room is the capitalised net operating income it supports, where μ converts room revenue to total revenue, m is the operating margin and c is the capitalisation rate:

$$V(t) = 365 A(t) O(t) \mu m / c \quad (5)$$

3.2 The entry trigger

A developer holding a site holds an option to build. Under the standard assumptions, with the value of a completed room following a geometric Brownian motion with volatility σ , a risk-adjusted discount rate r and a payout rate δ , the option is exercised when the value of the completed asset reaches a multiple of construction cost C given by:

$$V^* = \theta C, \theta = \beta / (\beta - 1) \quad (6)$$

$$\beta = 1/2 - (r - \delta)/\sigma^2 + \sqrt{\{[(r - \delta)/\sigma^2 - 1/2]^2 + 2r/\sigma^2\}} \quad (7)$$

Fig 4 and Table 4 report the multiple across the parameter range relevant to hotel development. It exceeds two for most of that range, which is far above what appraisal practice implies. Following Grenadier (1996, 2002) ^[18, 19] and Leahy (1993) ^[22] we therefore let competition erode the premium and write the effective trigger as:

$$\theta_{eff} = 1 + \lambda (\theta - 1), \lambda \in [0, 1] \quad (8)$$

Here λ is the share of the option premium that survives competitive pressure. At λ equal to one the marginal developer behaves as a monopolist with a secure claim to the site. At λ equal to zero the premium is fully dissipated and entry follows the net present value rule. The parameter is not a behavioural fudge: it is a reduced-form summary of site scarcity, the number of credible competitors and the security of development rights, all of which are in principle observable.

3.3 Preparation, starts and completions

The step that distinguishes this model from a standard stock-flow specification is the treatment of what happens while the trigger is unmet. Waiting is not passive. Sites are acquired, master plans are approved and designs are completed, so that projects accumulate in a state of readiness. Preparation responds to the net present value of a completed room, because a project that is not worth building at any horizon is not worth preparing:

$$I(t) = \psi K(t) \max\{0, V(t)/C - 1\} \quad (9)$$

Construction starts only when the effective trigger is met, at which point a fraction of the prepared stock breaks ground, subject to a physical constraint on how much can be started in one year:

$$S(t) = \min\{\zeta Q(t), s_{max} K(t), Q(t)\} \text{ if } V(t) \geq \theta_{eff} C; \quad S(t) = 0 \text{ otherwise} \quad (10)$$

$$Q(t+1) = (1 - \omega) Q(t) + I(t) - S(t) \quad (11)$$

Where ω is the rate at which approvals lapse or projects are abandoned. Rooms under construction and the standing stock then evolve mechanically:

$$\Pi(t) = S(t) + S(t-1) + \dots + S(t-\tau+1) \quad (12)$$

$$K(t+1) = (1 - \delta_K) K(t) + S(t-\tau) \quad (13)$$

The gate in equation (10) is what makes the trigger matter for the shape of the cycle rather than only for its average level. A higher trigger does not reduce the number of projects that are eventually built; it changes when they are built, and it allows a larger reservoir to accumulate before they are. Whether the net effect damps or amplifies the cycle is not obvious from inspection, and Section 5.3 answers it numerically.

3.4 The oversupply risk index

The index combines one contemporaneous, one forward-looking and one valuation component. The equilibrium stock is the number of rooms that would produce normal occupancy at current demand:

$$K^*(t) = D(t) / [365 O^*] \quad (14)$$

The occupancy gap measures the shortfall already realised:

$$G_1(t) = \max\{0, [O^* - O(t)] / O^*\} \quad (15)$$

The overhang component asks how many additional years of trend demand growth would be needed to absorb the stock that will exist when the current pipeline delivers, over and above the τ years it takes to deliver. A saturating transform keeps the component bounded without creating a mass point at the ceiling:

$$T(t) = \ln\{[K(t)(1 - \delta_K)^\tau + \Pi(t)] / [K^*(t)(1 + g)^\tau]\} / \ln(1 + g) \quad (16)$$

$$G_2(t) = 1 - \exp\{-\max[0, T(t)] / T_{ref}\} \quad (17)$$

The valuation component measures how far the value of a room has fallen below the cost of building one, which is the condition under which existing owners face impairment and new entry stops:

$$G_3(t) = \max\{0, [C - V(t)] / C\} \quad (18)$$

$$ORI(t) = w_1 G_1(t) + w_2 G_2(t) + w_3 G_3(t) \quad (19)$$

We use equal weights throughout. The choice is deliberate rather than optimal: an index intended for use by a planning authority should be transparent, and any weighting estimated on the history of one destination would not transfer to another. Section 5.7 reports how sensitive the results are to this choice.

3.5 Balanced growth and stability

Proposition 1. The model admits a balanced growth path on which occupancy equals O^* , the average daily rate is constant, and all stocks grow at rate g . Along that path, annual starts, the prepared stock and the value-to-cost ratio satisfy:

$$S / K = (g + \delta_K)(1 + g)^\tau, \quad Q / K = (S / K) / \zeta \quad (20)$$

$$V / C = \max\{\theta_{eff}, 1 + [(g + \omega)/\zeta + 1] (g + \delta_K)(1 + g)^\tau / \psi\} \quad (21)$$

The second term in (21) is the margin over construction cost needed to induce enough project preparation to replace depreciating rooms and to match demand growth. When the option trigger exceeds that margin, the trigger binds and the market operates at a permanently higher value-to-cost ratio. The proof is a direct substitution and is given in Appendix A.

Proposition 2. Linearising the system around the balanced growth path and writing $k(t)$ for the log deviation of the room stock yields a delay difference equation,

$$k(t+1) - k(t) = -\phi k(t-\tau) \quad (22)$$

Whose characteristic equation is $z^{\tau+1} (z - 1) + \phi = 0$. The composite gain ϕ is the product of the elasticity of starts with respect to the value gap and the elasticity of value with respect to the room stock. On the stability boundary the dominant root lies on the unit circle at an angle $\pi / (2\tau + 1)$, which gives:

$$\phi_{crit}(\tau) = 2 \sin[\pi / (2(2\tau + 1))], \text{ boundary cycle length} = 2(2\tau + 1) \text{ years} \quad (23)$$

Two things follow. The tolerance of the system for responsive construction falls sharply with the lag: a one-year lag permits a gain of one, a three-year lag permits only 0.445 and a five-year lag only 0.285. And the natural rhythm of the market is set by the lag alone. A three-year construction period implies a fourteen-year cycle at the boundary, which is close to the range reported for property markets generally, and which the model reproduces without any parameter chosen to that end. Fig 2 shows both relationships and Table 1 reports the root moduli at, above and below the boundary.




Fig 2: Left: the critical feedback gain falls steeply with the construction lag, so that a market with a longer delivery period tolerates far less responsive construction before it begins to oscillate. Right: the cycle length implied at the stability boundary, with the range commonly reported for property cycles shaded

Table 1: Stability of the linearised supply loop

τ (years)	Critical gain	Boundary cycle (years)	Root modulus at half the critical gain	Root modulus at twice the critical gain
1	1.000	6	0.707	1.414
2	0.618	10	0.820	1.225
3	0.445	14	0.869	1.155
4	0.347	18	0.897	1.118
5	0.285	22	0.915	1.096
6	0.241	26	0.928	1.080

Note. The dominant root modulus equals one exactly at the critical gain, by construction. A modulus above one implies an oscillation that grows rather than decays.

The practical importance of Proposition 2 is that it identifies the two quantities a planner can affect. Nothing in it depends on demand growth, on the level of construction costs or on the capitalisation rate. The stability of the market is a function of how quickly construction responds to profitability and of how long it takes to deliver, and of nothing else.

4. Calibration strategy and data requirements

4.1 Illustrative parameters

Table 2 lists the notation. Table 3 gives the parameter values used in every simulation reported below, together with the source from which each would be estimated in an empirical application. These values are illustrative. They were chosen to sit inside ranges reported in the literature and to produce a market whose balanced growth path is economically sensible, and they are not estimates of conditions in any Vietnamese city. Cells shown in highlight are the ones that must be replaced with estimated values before any result can be given an empirical interpretation.

Table 2: Notation

Symbol	Meaning	Unit
$K(t)$	Operating upscale rooms	rooms
$\Pi(t)$	Rooms under construction	rooms
$Q(t)$	Prepared projects not yet started	rooms
$S(t)$	Construction starts	rooms per year
$D(t)$	Room-night demand at the reference rate	room-nights
$O(t), O^*$	Occupancy and its normal level	share
$A(t)$	Average daily rate	currency
$V(t), C$	Value per room and development cost per room	currency
$\theta, \theta_{\text{eff}}$	Option trigger multiple and its competitive counterpart	ratio
λ	Share of the option premium surviving competition	share
τ	Construction lag	years
φ	Composite feedback gain of the supply loop	dimensionless
$\text{ORI}(t)$	Oversupply risk index	index, 0 to 1

Table 3: Illustrative parameter values and the source each would be estimated from

Symbol	Meaning	Value used
g	Trend growth of room-night demand	0.045
ρ, σ_D	Persistence and volatility of demand deviations	0.60, 0.05
O^*	Normal occupancy	0.68
O_{max}	Practical occupancy ceiling	0.92
α	Rate adjustment speed	0.50
η	Price elasticity of room-night demand	0.60
μ, m	Revenue multiple and operating margin	1.45, 0.26
c	Capitalisation rate	0.09
C	Development cost per room	105,000
τ	Construction lag	3
δ_K	Exit and obsolescence rate of rooms	0.02
ψ	Project preparation response	endogenous
ζ	Share of prepared stock started per year	0.50
ω	Lapse rate of approved projects	0.12
s_{max}	Ceiling on annual starts	0.16
r, δ, σ	Discount rate, payout rate, asset volatility	0.10, 0.06, 0.18
λ	Premium retention	0.35
T_{ref}	Reference absorption horizon in (17)	3 years
w_1, w_2, w_3	Index weights	1/3 each

Two parameters deserve comment. The capitalisation rate and the development cost per room have no effect whatsoever on any of the dynamic results reported below. This is not an artefact of the particular values chosen: the system is homogeneous of degree zero in the level of value, so only the ratio of value to cost enters the entry decision, and shifting both leaves every occupancy and index path unchanged. That property is useful in practice, because it means an empirical application does not need an accurate absolute construction cost, only an accurate ratio of what a completed room is worth to what it costs to build.

The preparation response ψ is not free. Given the other parameters, it is pinned down by the requirement that the balanced growth path exists, through equation (21). A market with a high effective trigger must have a

correspondingly low preparation response, since otherwise the reservoir of prepared projects would grow without bound. Imposing this restriction is what makes the comparison across entry regimes in Section 5.3 meaningful: each regime is compared at its own steady state rather than at an arbitrary common starting point.

4.2 What an empirical application requires

Applying the model to a specific coastal city requires four series and two cross-sectional estimates. Table 12 sets them out with their sources. The supply series is the most demanding item, because published room counts typically mix segments and often exclude projects that are licensed but not yet operating. The pipeline series is equally important and is rarely published in a form that distinguishes approved from started projects, which is precisely the distinction equation (10) turns on.

We note one constraint that will shape any Vietnamese application. Provincial tourism statistics report visitor arrivals and, in some provinces, room-nights sold, but the split between upscale and other accommodation is not consistently maintained, and the administrative reorganisation of provincial boundaries complicates the construction of a consistent series for some destinations. An empirical study would need to state explicitly how the destination market area is defined and to hold that definition constant across the sample period.

5. Model behaviour

5.1 Validation on the balanced growth path

With no shock and no demand noise, the model started on the path described by Proposition 1 stays there exactly. Occupancy remains at 0.680 for the full thirty-five year horizon, the rate is constant, and the risk index never exceeds 0.009, the small residual arising from the approximation of depreciation in equation (16). This is a useful check, because it establishes that everything reported below is a response to something and not an artefact of a badly set initial condition.

The path is a knife edge, however. Under the baseline parameters the composite gain is close to 0.5, slightly above the critical value of 0.445 implied by Proposition 2 for a three-year lag, so the balanced growth path is locally unstable. Any disturbance sets off an oscillation that does not die away. This is the central claim of the paper in its starkest form: the market does not need bad forecasts to overbuild, only a disturbance and a delay.

5.2 A permanent demand uplift

The first experiment applies a permanent eighteen per cent increase in room-night demand phased in from year five, of

the kind a new airport terminal, an expressway or a visa change might produce. The market response is shown in Fig 3.




Fig 3: Response to a permanent eighteen per cent demand uplift from year five. Top left: operating rooms, the equilibrium stock and rooms under construction. Top right: occupancy, against the trend-only path. Bottom left: rate and revenue per available room. Bottom right: the risk index and its three components

Occupancy rises immediately, because the stock cannot respond. The rate follows, the value of a room rises past the trigger, and the reservoir of prepared projects empties into construction. Nothing arrives for three years. When it does, it arrives in volume: the pipeline peaks at nearly half the operating stock, and the standing stock overshoots the equilibrium stock by year fifteen. Occupancy falls to 0.591, twelve percentage points below normal, and revenue per available room falls by 37 per cent from its peak. The market then repeats the pattern at lower amplitude, with a second trough around year thirty-two.

The decomposition in the bottom right panel is instructive for practical monitoring. The overhang component moves first and carries most of the index; the occupancy gap follows several years later; the valuation component barely moves at all under this scenario, because the rate never falls far enough to push the value of a room below its construction cost. An authority monitoring occupancy alone would have seen nothing untoward until the rooms were already under construction.

5.3 Competitive erosion of the option premium

The textbook trigger for these parameters is 2.17, meaning a completed room should be worth more than twice its construction cost before a protected developer would build. Table 4 shows how that multiple varies with volatility and the discount rate, and Fig 4 plots it. Observed behaviour in competitive development markets sits far below this, which is exactly what the pre-emption literature predicts.




Fig 4: The entry trigger multiple implied by the standard real options solution, as a function of asset value volatility, for three discount rates, with a payout rate of six per cent. The horizontal line is the net present value rule

Table 4: Entry trigger multiple θ for combinations of volatility and discount rate

Volatility σ	$r = 0.08$	$r = 0.10$	$r = 0.12$	β at $r = 0.10$
0.10	1.564	1.848	2.155	2.179
0.15	1.766	2.035	2.329	1.966
0.18	1.903	2.168	2.456	1.856
0.20	2.000	2.264	2.549	1.791
0.25	2.266	2.528	2.809	1.654
0.30	2.563	2.827	3.106	1.547

Note. Payout rate fixed at 0.06. The value 2.168 corresponds to the illustrative baseline used throughout.

Table 5 and Fig 5 trace what happens as the surviving share of the premium is varied from zero to one, holding everything else constant and comparing each regime at its own balanced growth path. The relationship is monotone and strong. With the premium fully dissipated, occupancy falls to 0.551 at the trough and the risk index peaks at 0.359. With the premium fully retained, the trough is 0.633 and the index peaks at 0.126. Erosion of the option premium accounts for a swing of roughly eight percentage points in trough occupancy, which is larger than the effect of any other parameter we vary.




Fig 5: Left: occupancy paths after the demand uplift, for three values of the share of the option premium surviving competition. Right: cycle severity across the full range of that share

Table 5: Cycle severity as the option premium is competitively eroded

λ	θ_{eff}	Trough occupancy	Occupancy amplitude	Peak risk index	RevPAR drawdown
0.00	1.000	0.551	0.284	0.359	0.252
0.10	1.117	0.551	0.296	0.336	0.347
0.20	1.234	0.562	0.276	0.311	0.412
0.30	1.350	0.595	0.235	0.258	0.339
0.35	1.409	0.591	0.240	0.258	0.368
0.40	1.467	0.608	0.207	0.222	0.272
0.50	1.584	0.616	0.203	0.196	0.271
0.60	1.701	0.618	0.202	0.184	0.289
0.80	1.934	0.621	0.204	0.151	0.243
1.00	2.168	0.633	0.172	0.126	0.191

The interpretation matters for policy. The literature on overbuilding tends to read it as a failure of information or of judgement, which invites remedies built around better forecasting and better disclosure. The model says something different. A developer who correctly values the option to wait, and who would exercise that option if allowed to, is forced by the threat of pre-emption to build earlier than is individually optimal. The result is a market that behaves as though nobody valued waiting at all. The problem is one of coordination, not of information, and instruments aimed at information will therefore underperform.

5.4 The construction lag

Proposition 2 says the lag governs stability. Table 6 and Fig 6 confirm that the nonlinear model behaves as the linearised analysis predicts. At a one-year lag the critical gain is 1.0, comfortably above the baseline gain of roughly 0.5, and cycle is shallow, with a trough of 0.622. At a two-year lag the critical gain is 0.618, still above 0.5, and the trough falls only to 0.611. At three years the critical gain drops below the operating gain and the market crosses into the oscillatory region: the trough falls to 0.591 and the pipeline peaks at 49 per cent of the standing stock. Beyond that, each additional year of lag makes matters worse, and at a five-year lag the pipeline peaks at 85 per cent of the operating stock.




Fig 6: Left: occupancy paths after the demand uplift for construction lags of two to five years. Right: cycle severity by lag. The transition between the second and third bars corresponds to the crossing of the analytical stability boundary

Table 6: Cycle severity by construction lag

τ (years)	Critical gain	Boundary cycle (years)	Trough occupancy	Peak risk index	Peak pipeline / stock
1	1.000	6	0.622	0.145	0.14
2	0.618	10	0.611	0.189	0.31
3	0.445	14	0.591	0.258	0.49
4	0.347	18	0.596	0.274	0.61
5	0.285	22	0.591	0.292	0.85

Note. The operating gain under the baseline parameters is approximately 0.5. Rows in which the critical gain falls below that value lie in the oscillatory region identified by Proposition 2.

5.5 A transitory collapse

The third experiment is a sudden fall in demand of 45 per cent in year twelve, followed by a geometric recovery, which is the shape observed in the international tourism collapse of 2020 and 2021 as documented by Gössling, Scott and Hall (2021)^[16] and Škare, Soriano and Porada-Rochón (2021)^[31]. Fig 7 shows the result.




Fig 7: A transitory demand collapse in year twelve with geometric recovery. Left: occupancy. Right: the risk index and its forward-looking pipeline component

Occupancy falls to 0.374 and the index jumps to 0.555, well into what we classify as the severe band. The index behaves as it should here: the collapse is registered immediately by the occupancy and valuation components rather than by the pipeline component, which is exactly the right diagnosis, since the problem in a collapse is demand rather than commitments. The more interesting feature is what happens afterwards. Construction stops entirely for several years, the stock depreciates, and the recovery in demand meets a shrunken supply, which drives occupancy to 0.877 and starts the whole cycle again. The echo is visible in the index from around year twenty-six. Table 7 sets the three scenarios side by side.

Table 7: Scenario summary

Scenario	Trough occupancy	Occupancy amplitude	Peak risk index	Peak pipeline / stock
S0. Trend growth only	0.680	0.000	0.009	0.20
S1. Permanent demand uplift of 18 per cent	0.591	0.240	0.258	0.49
S2. Transitory collapse of 45 per cent	0.374	0.503	0.555	0.41

Note. All scenarios use the baseline parameters of Table 3 with a premium retention share of 0.35 and a three-year construction lag.

5.6 Stochastic demand

The deterministic experiments isolate mechanisms but say nothing about how often trouble occurs. We therefore two thousand paths with the persistent demand noise of equation (2) and no scenario shock. Fig 8 and Table 8 report the distribution.




Fig 8: Two thousand paths under persistent demand noise. Left: the distribution of occupancy over time. Right: the distribution of the highest risk index value reached on each path

Ordinary demand variation, with no large shock at all, is enough to produce an oversupply episode on almost every path. Occupancy falls below 0.60 at some point on 96 per cent of paths, and below 0.55 on 46 per cent. Revenue per available room falls by an average of 34 per cent from peak to trough, and by more than half on the worst five per cent

of paths. The median peak of the risk index is 0.31 and nine paths in ten reach the elevated band.

The index also leads. On 70 per cent of paths it enters the elevated band before the occupancy trough, with a median lead of four years and a mean of eight. Four years is roughly the time needed to stop a project that has been approved but not started, which is the decision the index is designed to inform. Table 9 sets out the bands we propose for classifying the index. They are anchored on the distribution just described and would have to be re-anchored on a calibrated market.

Table 8: Distribution of outcomes across two thousand paths

Statistic	Premium retention 0	Baseline (0.35)	Premium retention 1
Mean trough occupancy	0.528	0.554	0.575
Probability occupancy falls below 0.60	97.3%	95.8%	86.3%
Probability the index reaches the elevated band	99.0%	88.3%	53.3%
Mean RevPAR drawdown	36.8%	34.0%	28.8%

Table 9: Proposed bands for the oversupply risk index

Range	Classification	Interpretation and suggested response
below 0.15	Balanced	Pipeline consistent with demand growth. Routine monitoring.
0.15 to 0.25	Watch	Commitments running ahead of absorption. Review approvals not yet started.
0.25 to 0.35	Elevated	Absorption of the pipeline requires several additional years of growth. Consider phasing new approvals.
0.35 and above	Severe	Either a realised occupancy shortfall or an overhang that trend growth cannot absorb within the planning horizon.

5.7 Sensitivity

Table 10 varies each parameter one at a time across a plausible interval and reports the effect on trough occupancy. The ranking is stable and has a clear message. The two parameters that matter most are the share of the option premium that survives competition and the price elasticity of demand. The first governs how synchronised entry is; the second governs how much of a supply overhang can be absorbed by discounting rather than by empty rooms. The construction lag and the trend growth rate follow. The capitalisation rate and the development cost per room have no effect at all, for the homogeneity reason given in Section 4.1.

Table 10: One-at-a-time sensitivity of trough occupancy

Parameter	Low value	High value	Trough at low	Trough at high
Premium retention λ	0.20	0.50	0.562	0.616
Price elasticity of demand	0.40	0.80	0.572	0.612
Trend growth of demand	0.030	0.060	0.614	0.590
Construction lag (years)	2	4	0.611	0.596
Lapse rate of approvals	0.06	0.18	0.580	0.594
Share of reservoir started per year	0.35	0.65	0.599	0.590
Rate adjustment speed	0.35	0.65	0.600	0.601
Capitalisation rate	0.075	0.105	0.591	0.591
Development cost per room	90,000	120,000	0.591	0.591

Note. Baseline trough occupancy is 0.591. All runs use the permanent demand uplift of Section 5.2. Parameters are ordered by the width of the induced range.

We also varied the index weights. Shifting weight from the occupancy gap to the overhang component raises the index earlier in the cycle and lengthens the lead time, at the cost of more frequent signals that are not followed by a trough. The equal weighting used throughout is a compromise rather than an optimum, and an authority that cares more about lead time than about precision should weight the overhang component more heavily.

6. Policy experiments

Four instruments are compared, all applied to the permanent demand uplift of Section 5.2. A cap limits construction starts in any year to a fixed share of the operating stock. A disclosure regime requires the pipeline to be published, which we model as developers netting the excess pipeline out of their valuation, so that a large overhang reduces the perceived value of a new project. A lag reduction shortens the construction period from three years to two, standing in for faster permitting and modular construction. The fourth is the cap and disclosure combined. Fig 9 and Table 11 report the outcomes.




Fig 9: Policy experiments under the permanent demand uplift. Left: occupancy paths. Right: trough occupancy and the peak risk index under each instrument

Table 11: Policy experiments

Instrument	Trough occupancy	Occupancy amplitude	Peak risk index	Peak pipeline / stock
No intervention	0.591	0.240	0.258	0.49
Cap on starts at 8 per cent of stock	0.667	0.099	0.009	0.22
Cap on starts at 5 per cent of stock	0.680	0.122	0.009	0.20
Pipeline disclosure	0.620	0.178	0.151	0.42
Construction lag shortened to two years	0.611	0.185	0.189	0.31
Cap at 8 per cent plus disclosure	0.667	0.119	0.009	0.23

Note. All instruments are applied from the start of the simulation. Amplitude is the difference between the highest and lowest occupancy over the horizon.

The cap dominates. Limiting starts to eight per cent of the operating stock in any year raises trough occupancy from 0.591 to 0.667 and effectively removes the episode: the index never leaves the balanced band, and the pipeline never exceeds 22 per cent of the standing stock. The reason is visible in Proposition 2. A cap does not change the trigger or the lag; it truncates the response of starts to profitability, which is precisely the gain that has to fall below the critical value for the system to be stable.

Two caveats keep this from being a simple recommendation. The tighter cap of five per cent produces a higher trough but

a larger amplitude, because the market spends longer undersupplied and occupancy overshoots on the upside. A cap that binds too tightly converts an oversupply problem into a scarcity problem, with consequences for rates, for visitor numbers and for the destinations that capture the diverted demand. And a cap allocates scarce permission, which raises questions of rationing and of rent capture that fall outside this model but not outside the problem.

Disclosure and lag reduction both help, and both help less than the cap. Disclosure lifts the trough by roughly three percentage points. Shortening the lag lifts it by two and also halves the peak pipeline, since fewer years of starts are in the ground at any moment. Combining the cap with disclosure adds nothing measurable to the cap alone, which is consistent with the reading of Section 5.3: once the quantity of starts is constrained, the information available to developers is no longer the binding constraint.

7. Discussion

The model reframes overbuilding in coastal resort markets as a control problem rather than an information problem. The three ingredients that generate it, a delivery lag measured in years, a threshold entry rule, and a stock of prepared projects released in a burst, are all observable, and none requires any participant to be wrong about anything. That has consequences for how the problem is discussed and for what is done about it.

The first consequence concerns the diagnosis. When a destination finds itself with too many rooms, the post-mortem usually identifies over-optimistic feasibility studies. Our results suggest that even feasibility studies that were correct at the time would have produced the same aggregate outcome, because each was correct about the market as it stood and none could condition on what every other developer was about to do. The failure is in the absence of a mechanism that aggregates individual decisions, not in the decisions themselves.

The second concerns the instrument. Most destination-level responses to overbuilding operate through information: better market studies, published pipeline data, investor education. The model says these help at the margin and that a quantity restriction helps far more. It also says why, in a way that is testable: the binding quantity in equation (10) is the flow of starts, and the stability condition in (23) is a condition on the responsiveness of that flow.

The third concerns timing. Because the index leads the trough by a median of four years in the stochastic simulations, and because the decision it informs is whether to allow an approved project to break ground, the monitoring and the instrument fit together. An authority that publishes the index annually and ties the start cap to its band has a rule that is simple, transparent and aligned with the mechanism the model identifies. Whether such a rule survives contact with the political economy of provincial development approval is a question this paper cannot answer.

A fourth point is comparative. The stability boundary in (23) does not contain a single tourism-specific parameter, so it applies equally to offices, apartments and logistics space. What makes accommodation markets more exposed is the speed with which the space market clears. A hotel reprices nightly, so a supply shock passes into revenue and into asset value within a year, which places the composite gain higher than it would be for a market with multi-year leases. The

same lag, in other words, is more dangerous in hotels than in offices.

8. Limitations

The most important limitation is that the model has not been estimated. The parameters in Table 3 are illustrative, and while they are chosen from defensible ranges, the quantitative results should be read as statements about the behaviour of the model, not about any actual destination. Section 4.2 and Table 12 set out what an empirical application requires, and assembling those series is the obvious next step.

The model treats the destination as a single homogeneous market. Real coastal cities contain submarkets that are imperfect substitutes, separated by beachfront access, by distance from the airport and by brand affiliation, and a pipeline concentrated in one submarket does not compete evenly with the whole stock. Adding that heterogeneity would raise the effective absorption capacity of the destination and would probably lower the amplitude of the cycle.

Developers are represented by a single decision rule with a common trigger. In practice, entry decisions differ across investor types, and an international operator financed offshore faces a different cost of capital, a different tolerance for a slow ramp-up and a different exit option from a domestic developer. Heterogeneous triggers would smooth the gate in equation (10) and would reduce the synchronisation that drives our main result, so our estimate of the cost of competitive premium erosion should be read as an upper bound.

Finally, the demand side is exogenous. Room supply in a resort destination is not a pure response to visitor numbers: new capacity brings marketing, new routes and new attractions, and can create part of the demand it goes on to serve. Incorporating that feedback would make oversupply less persistent than the model implies, though at the cost of a considerably more complex system and of parameters that are harder still to identify.

9. Conclusion

Oversupply in coastal hotel markets is usually treated as a forecasting failure. This paper has argued, and shown within a formal model, that it is better understood as the predictable behaviour of a delayed feedback system whose participants are individually rational. Three results carry the argument. The balanced growth path of the market is locally unstable once construction responds to profitability more strongly than a threshold that falls steeply with the construction lag, and that threshold has a closed form. The natural rhythm of the resulting cycle is set by the lag alone, and a three-year delivery period reproduces the observed decade-and-a-half rhythm of property cycles without any tuning. And the competitive erosion of the option to wait, which makes each developer build earlier than would be individually optimal, is the single most powerful driver of how deep the trough goes.

For a destination authority the practical implications are narrow and specific. Monitor the pipeline as a ratio to the standing stock and to the demand it will meet on delivery, not as a raw count of rooms. Expect the warning to arrive several years before the occupancy data confirm it. And if an instrument is to be used, constrain the quantity of starts rather than attempting to improve the information on which

the starts are based.

The framework is deliberately small, and that is what makes it usable. Every equation is written in quantities a provincial planning office either holds or could collect within a year, and the next step is to collect them for one coastal city and put the model to the test.

10. References

1. Assaf AG, Josiassen A. Frontier analysis: A state-of-the-art review and meta-analysis. *Journal of Travel Research*. 2016; 55(5):612-627.
2. Bandalouski AM, Kovalyov MY, Pesch E, Tarim SA. An overview of revenue management and dynamic pricing models in hotel business. *RAIRO Operations Research*. 2018; 52(1):119-141.
3. Bar-Ilan A, Strange WC. Investment lags. *American Economic Review*. 1996; 86(3):610-622.
4. Barros CP. Measuring efficiency in the hotel sector. *Annals of Tourism Research*. 2005; 32(2):456-477.
5. Borio C, Drehmann M. Assessing the risk of banking crises, revisited. *BIS Quarterly Review*, March 2009, 29-46.
6. Borio C, Lowe P. Asset prices, financial and monetary stability: Exploring the nexus. *BIS Working Paper 114*. Bank for International Settlements, 2002.
7. Butler RW. The concept of a tourist area cycle of evolution: Implications for management of resources. *Canadian Geographer*. 1980; 24(1):5-12.
8. Capozza DR, Helsley RW. The stochastic city. *Journal of Urban Economics*. 1990; 28(2):187-203.
9. Coccossis H, Mexa A. (Eds.). *The challenge of tourism carrying capacity assessment: Theory and practice*. Ashgate, 2004.
10. DiPasquale D, Wheaton WC. The markets for real estate assets and space: A conceptual framework. *Journal of the American Real Estate and Urban Economics Association*. 1992; 20(1):181-198.
11. Dixit AK, Pindyck RS. *Investment under uncertainty*. Princeton University Press, 1994.
12. Drehmann M, Juselius M. Evaluating early warning indicators of banking crises: Satisfying policy requirements. *International Journal of Forecasting*. 2014; 30(3):759-780.
13. Enz CA, Canina L, Lomanno M. Competitive pricing decisions in uncertain times. *Cornell Hospitality Quarterly*. 2009; 50(3):325-341.
14. Forrester JW. *Industrial dynamics*. MIT Press, 1961.
15. Getz D. Tourism planning and destination life cycle. *Annals of Tourism Research*. 1992; 19(4):752-770.
16. Gössling S, Scott D, Hall CM. Pandemics, tourism and global change: A rapid assessment of COVID-19. *Journal of Sustainable Tourism*. 2021; 29(1):1-20.
17. Grenadier SR. Local and national determinants of office vacancies. *Journal of Urban Economics*. 1995; 37(1):57-71.
18. Grenadier SR. The strategic exercise of options: Development cascades and overbuilding in real estate markets. *Journal of Finance*. 1996; 51(5):1653-1679.
19. Grenadier SR. Option exercise games: An application to the equilibrium investment strategies of firms. *Review of Financial Studies*. 2002; 15(3):691-721.
20. Kaminsky G, Lizondo S, Reinhart CM. Leading indicators of currency crises. *IMF Staff Papers*. 1998; 45(1):1-48.

21. Kydland FE, Prescott EC. Time to build and aggregate fluctuations. *Econometrica*. 1982; 50(6):1345-1370.
22. Leahy JV. Investment in competitive equilibrium: The optimality of myopic behavior. *Quarterly Journal of Economics*. 1993; 108(4):1105-1133.
23. Liu H, Liu Y, Wang Y, Pan C. Hot topics and emerging trends in tourism forecasting research: A scientometric review. *Tourism Economics*. 2019; 25(3):448-468.
24. Majd S, Pindyck RS. Time to build, option value, and investment decisions. *Journal of Financial Economics*. 1987; 18(1):7-27.
25. McDonald R, Siegel D. The value of waiting to invest. *Quarterly Journal of Economics*. 1986; 101(4):707-727.
26. Pindyck RS. Irreversible investment, capacity choice, and the value of the firm. *American Economic Review*. 1988; 78(5):969-985.
27. Qiu RTR, Wu DC, Dropsy V, Petit S, Pratt S, Ohe Y. Visitor arrivals forecasts amid COVID-19: A perspective from the Asia and Pacific team. *Annals of Tourism Research*. 2021; 88:103155.
28. Saveriades A. Establishing the social tourism carrying capacity for the tourist resorts of the east coast of the Republic of Cyprus. *Tourism Management*. 2000; 21(2):147-156.
29. Sedarati P, Santos S, Pintassilgo P. System dynamics in tourism planning and development. *Tourism Planning and Development*. 2019; 16(3):256-280.
30. Sigala M. Tourism and COVID-19: Impacts and implications for advancing and resetting industry and research. *Journal of Business Research*. 2020; 117:312-321.
31. Škare M, Soriano DR, Porada-Rochón M. Impact of COVID-19 on the travel and tourism industry. *Technological Forecasting and Social Change*. 2021; 163:120469.
32. Smeral E. Tourism forecasting performance considering the instability of demand elasticities. *Journal of Travel Research*. 2017; 56(7):913-926.
33. Song H, Li G. Tourism demand modelling and forecasting: A review of recent research. *Tourism Management*. 2008; 29(2):203-220.
34. Song H, Qiu RTR, Park J. A review of research on tourism demand forecasting. *Annals of Tourism Research*. 2019; 75:338-362.
35. Sterman JD. *Business dynamics: Systems thinking and modeling for a complex world*. McGraw-Hill, 2000.
36. Titman S. Urban land prices under uncertainty. *American Economic Review*. 1985; 75(3):505-514.
37. Wheaton WC. The cyclic behavior of the national office market. *AREUEA Journal*. 1987; 15(4):281-299.
38. Wheaton WC, Rossoff L. The cyclic behavior of the U.S. lodging industry. *Real Estate Economics*. 1998; 26(1):67-82.