Journal of Industrial and Management Optimization

https://www.aimspress.com/journal/jimo

JIMO, 22(10): 4853–4900.
DOI:10.3934/jimo.2026168
Received: 21 June 2026
Revised: 21 August 2026
Accepted: 03 September 2026
Published: 10 September 2026

Research article

Bidirectional calibration-robust booking control under heterogeneous cancellation risk

Jinho Cha[1] **and Sook Young Lim**[2,*]

[1] Department of Computer Science, Gwinnett Technical College, Lawrenceville, GA, USA

[2] Department of Airline Service, Kyungwoon University, Gumi, Republic of Korea

* **Correspondence:** Email: synangja@ikw.ac.kr.

Abstract: Reservation-based capacity systems often rely on predicted cancellation or show-up probabilities when deciding how much demand to accept before service realization. Existing hotel booking and overbooking models typically treat these probabilities as fixed inputs or merge different forecast errors into a single uncertainty source. Less is known about how a capacity-constrained decision-maker should protect a segment-level acceptance portfolio when heterogeneous show-up probabilities are imperfectly calibrated and aggregate residual load remains distributionally ambiguous. This paper develops a calibration-robust optimization framework for this problem. Segment-level probability miscalibration is modeled by a weighted budgeted uncertainty set, while ambiguity in the aggregate residual load law is modeled by a 1-Wasserstein ambiguity set. The robust booking control model combines asymmetric vacancy and overflow costs with a distributionally robust conditional value-at-risk (CVaR) service-risk constraint. The main theoretical contribution is the bidirectional calibration principle: The locally adverse probability perturbation is segment specific and reverses when the marginal value of capacity crosses a segment-specific value-to-load threshold. We derive finite reformulations and a certified scenario-and-piece cutting-plane algorithm with finite convergence for the finite piecewise-linear model. Controlled hotel-booking experiments check the direction-reversal mechanism, quantify the separate and joint effects of the two uncertainty layers, and evaluate scalability. A public hotel booking data set is used only as a real-world-calibrated illustration of segment heterogeneity and calibration bands. The results show that separating calibration risk from aggregate load ambiguity changes the portfolio composition, capacity protection, and the auditable price of robustness.

Keywords: robust booking control; calibration uncertainty; distributionally robust optimization; Wasserstein ambiguity; hotel revenue management
Mathematics Subject Classification: 90C15, 90C17, 90B50, 91B32

1. Introduction

Reservation-based capacity systems commit scarce, perishable resources before service realization. Hotels accept room reservations, airlines sell seats, clinics schedule appointments, and rental platforms promise finite inventory even though some commitments later cancel or fail to materialize. The operational trade-off is inherently asymmetric: Insufficient acceptance creates unused capacity that cannot be recovered after the service date, whereas excessive acceptance creates overflow, denial, displacement, or service recovery costs. Classical revenue management models establish how perishable capacity should be controlled under stochastic demand and advance commitments [1, 2]. Hotel-specific studies extend this logic to room allocation and cancellation-sensitive booking decisions [3, 4]. Overbooking models further quantify the trade-off between unused capacity and denied service [5, 6]. In hotel booking control, however, this trade-off is also segment dependent. Group bookings may have lower show-up rates but larger room-night loads, direct bookings may have higher show-up rates and narrower calibration bands, and online travel agency (OTA) bookings may combine large volume with greater cancellation uncertainty. These differences make the problem a capacity allocation problem under uncertain probability inputs, not only a cancellation forecasting problem.

The optimization challenge becomes sharper when booking decisions rely on heterogeneous, model-based show-up probabilities. Once probability estimates enter a booking-control model, two distinct forms of uncertainty arise. First, nominal segment-level show-up probabilities may be imperfectly calibrated because of finite samples, temporal drift, segment imbalance, or model misspecification. Empirical hospitality research documents substantial variation in cancellation behavior across lead time, distribution channel, contract terms, pricing, customer history, and market segment [7, 8]. Interpretable and profit-aware studies confirm that the numerical quality of reservation-level probabilities matters for operational use [9–11]. Prediction assisted booking models increasingly exploit such heterogeneity [12,13], while decision-based model selection research shows that predictive fit and downstream decision quality need not coincide [14]. However, these studies generally use the resulting probabilities as fixed decision inputs. The unresolved decision problem is not simply whether cancellations can be predicted. It is how a capacity-constrained operator should choose segment-level booking limits when the predicted probabilities themselves may be miscalibrated and when the aggregate realized load remains distributionally ambiguous. This distinction motivates a booking control model that separates segment-level calibration uncertainty from ambiguity in the conditional aggregate-load law.

Recent robust and distributionally robust optimization (DRO) research provides the methodological foundation for protecting decisions against parameter error and ambiguity in probability laws. Budgeted uncertainty sets control the price of robustness while preserving tractability [15]. Moment-based ambiguity sets protect against incomplete distributional information [16], whereas Wasserstein ambiguity sets retain geometric information about distributional perturbations and admit data-driven guarantees [17, 18]. Two-stage and general Wasserstein formulations establish tractable robust counterparts under suitable regularity conditions [19,20]. Recent work further clarifies variation– regularization effects [21], general transport duality [22], and the common convex-analytic structure of robust and distributionally robust optimization [23]. Continuity and optimizer existence results provide additional support for ambiguity radius sensitivity analysis [24]. Prescriptive analytics research has also begun to distinguish uncertainty in predictive parameters from ambiguity in residual distributions [25, 26]. In the present model, these tools play different roles: The calibration set perturbs segment-level show-up inputs, whereas the

Wasserstein set protects the conditional aggregate load law. Keeping these layers separate is essential because they affect different economic primitives and can lead to different booking portfolios.

The central observation of this paper is that the economically adverse direction of probability error is not uniform across segments. A downward perturbation in a segment's show-up probability reduces realized contractual value but also relieves capacity congestion. An upward perturbation increases realized value but can intensify overflow and service recovery exposure. Which direction is adverse therefore depends on the segment's incremental value relative to its capacity burden and on the endogenous marginal price of capacity. We call this phenomenon the *bidirectional calibration principle*. The adverse endpoint reverses when the marginal capacity price crosses a segment specific value-to-load threshold, generating a phase-transition structure that cannot be obtained from one-directional stress tests or aggregate rate perturbations. Thus, the contribution is not merely to acknowledge heterogeneous cancellation risk, but to show that the heterogeneous value-load structure makes the worst-case calibration direction itself endogenous.

To formalize this idea, the paper develops a multidimensional booking control model in which the decision-maker chooses accepted quantities across heterogeneous reservation segments rather than a scalar overbooking buffer. Segment-level calibration uncertainty is represented by a weighted budgeted set around nominal show-up probabilities, following the tractable robustness logic of [15]. Conditional aggregate load is represented by a standardized lower-truncated skew-normal reference law, and ambiguity around that law is captured by a 1-Wasserstein neighborhood. The objective maximizes the worst-case contractual value net of asymmetric overflow and unused capacity losses, while a distributionally robust conditional value-at-risk (CVaR) constraint controls service exposure using the convex risk representation of [27]. The Wasserstein component is supported by modern strong duality and finite reformulation results [17, 20, 22, 23]. Figure 1 summarizes the modeling flow while avoiding notation that is introduced only later in the formulation.

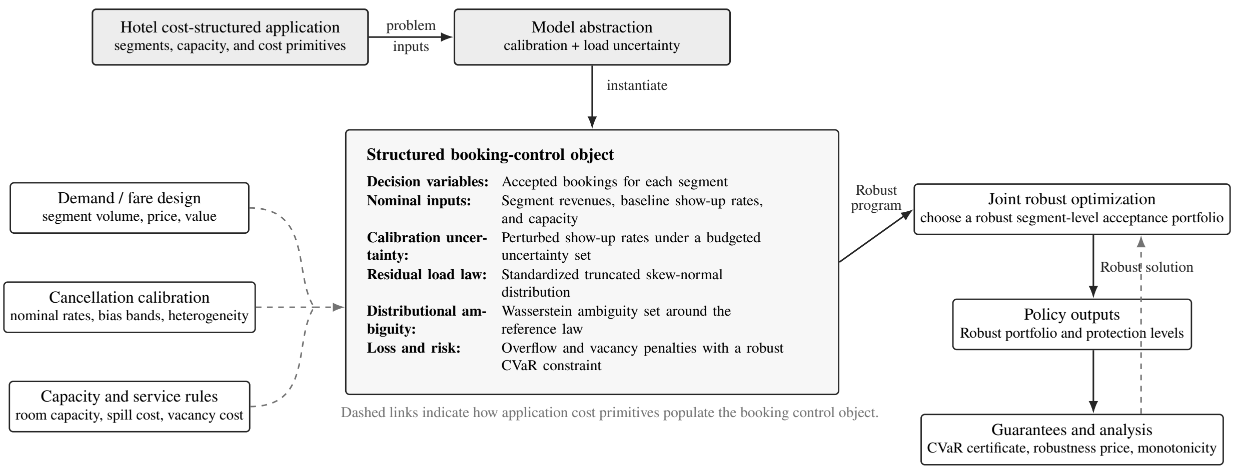

Figure 1. Architecture of the proposed calibration robust booking control framework. The figure separates hotel operating inputs, model abstraction, structured booking control objects, robust optimization, policy outputs, and analytical guarantees. Dashed links indicate how application cost primitives and operating assumptions populate the booking control object.

This paper makes four contributions.

1. It formulates a robust booking control model that separates segment-level calibration uncertainty from ambiguity in the aggregate residual load distribution. The formulation combines a weighted budgeted calibration set, a 1-Wasserstein ambiguity set, asymmetric vacancy and overflow costs, and a distributionally robust conditional value-at-risk service risk constraint.

2. It establishes the bidirectional calibration principle. The adverse probability perturbation is not uniformly upward or downward; its direction is segment-dependent and reverses when the marginal value of capacity crosses a value-to-load threshold.

3. It derives a finite piecewise-linear robust counterpart and a certified scenario-and-piece cutting-plane algorithm with finite convergence for the model solved in computation. The algorithm returns a robust segment portfolio together with active scenario, residual, and finite gap certificates.

4. It uses a hotel cost-structured application to interpret how vacancy cost, overflow cost, service-risk tolerance, calibration ambiguity, and distributional ambiguity jointly reshape the acceptance portfolio and the price of robustness.

The hotel setting provides a cost-structured application in which vacancy losses, overflow penalties, service risk tolerance, and capacity protection trade offs have direct economic interpretations. Controlled hotel booking instances are complemented by a public data illustration that grounds segment shares, cancellation levels, revenue proxies, stay load proxies, and calibration bands. The public records are used as a real-world-calibrated illustration of the model inputs and segment heterogeneity.

The remainder of the paper is organized as follows. Section 2 positions the model relative to reservation-based capacity control, robust and distributionally robust optimization, heterogeneous cancellation risk, and probability calibration. Section 3 formulates the reservation capacity problem. Section 4 develops the structural theory. Section 5 derives tractable reformulations and the finite convergence algorithm. Section 6 specifies the computational experiments and real-world-calibrated illustration. Section 7 reports numerical results, including uncertainty set sensitivity, algorithmic efficiency, and cost policy interpretation. Section 8 discusses the theoretical and policy implications, and Section 9 concludes. Proofs are collected in Appendix A.

2. Related literature and positioning

2.1. *Hotel booking control, cancellation heterogeneity, and prediction-assisted decisions*

Booking control under cancellations and no-shows is a foundational problem in revenue management and operations research. Classical work studies perishable inventory, stochastic arrivals, booking limits, bid price controls, and overbooking decisions in transportation and service systems [1, 2, 5, 28]. Hotel revenue management research extends these ideas to room allocation, rolling horizons, multi-night capacity interactions, and hotel-specific demand structure [3, 4, 29–31]. Models that combine cancellations, customer choice, and overbooking further demonstrate that cancellation behavior changes both the availability and the economic value of accepted demand [6, 32]. Recent hotel revenue management research further emphasizes that overbooking performance depends on forecast accuracy, segmentation, overselling costs, and the method used to set overbooking levels [33].

This literature establishes the spoilage–overflow trade-off that underlies our objective. It also provides the natural optimization language for the present paper: Capacity is finite and perishable, accepted commitments create future random load, and the decision-maker must balance the unused capacity cost against the service disruption cost. It also clarifies why hotel booking control is naturally segment dependent. Empirical hospitality research documents substantial variation in cancellation behavior across lead time, distribution channel, contract terms, pricing, customer history, stay characteristics, and market segment [7–9, 34–36]. Recent studies emphasize interpretable, adaptive, dynamic, or profit-sensitive cancellation prediction rather than discrimination alone [10, 11, 37, 38]. This evidence justifies segment-specific nominal show-up probabilities in our model.

The closest hotel-specific optimization studies combine overbooking with allocation or prediction-assisted control. Pimentel et al. optimized simultaneous hotel overbooking and market segment allocation in a synthetic environment, while Zhai et al. fed individualized no-show estimates into hotel overbooking decisions [12, 13]. Decision-based model-selection research further shows that predictive fit and downstream decision quality need not coincide [14]. These studies motivate heterogeneous booking decisions, but the estimated probabilities remain fixed inputs. Our model instead treats nominal probabilities as uncertain optimization primitives and derives segment-specific adverse endpoints, acceptance priorities, and cost consequences under robust protection.

Thus, the relevant question is not only whether cancellation risk can be predicted. It is how a capacity allocation decision should be protected when heterogeneous show-up probabilities may be miscalibrated and when the aggregate realized load remains distributionally ambiguous. Most booking control models represent cancellations or no-shows through known class probabilities, historical averages, exogenous scenarios, or jointly specified stochastic processes. Risk-averse and robust revenue management studies relax some of these assumptions and protect decisions against forecast error or uncertain demand distributions [39, 40]. However, they generally do not isolate segment-level probability miscalibration from ambiguity in the conditional aggregate load law, nor do they characterize the direction of a worst-case probability perturbation as an endogenous function of capacity scarcity.

2.2. *Robust and distributionally robust optimization under forecast uncertainty*

Robust optimization protects decisions against uncertain parameters by optimizing over prescribed sets. The budgeted uncertainty framework of Bertsimas and Sim is especially useful because it allows only a limited number or weighted magnitude of coefficients to deviate simultaneously, producing an interpretable robustness budget and tractable counterparts [15]. We adapt this logic to calibration errors around nominal segment show-up probabilities. Unlike a standard one-sided coefficient stress, however, the sign of the damaging deviation is endogenous and segment–specific.

Distributionally robust optimization protects decisions against uncertainty in the probability law itself. Moment ambiguity [16], Wasserstein ambiguity [17,18,20], and two-stage Wasserstein formulations [19] provide tractable ways to control distributional misspecification. Applications to capacity-constrained service systems demonstrate the practical value of such protection [41]. Reviews have synthesized the broader modeling and computational landscape [40].

Several recent theoretical advances are particularly relevant. Gao et al. interpreted Wasserstein DRO as a regularization of the variation of the loss function, clarifying how the ambiguity radius penalizes distributional sensitivity [21]. Zhang et al. provided a concise strong duality argument for Wasserstein DRO under general transport costs and measurable losses [22]. Zhen et al. developed a unified convex-analytic principle connecting robust and distributionally robust optimization and derive finite convex reformulations for broad optimal-transport ambiguity sets [23]. Fan et al. established the existence and continuity properties for the Wasserstein distributional model risk as the radius changes [24]. These results support our dual reformulation, sensitivity interpretation, monotonicity analysis, and uncertainty-set sensitivity experiments.

These tools have also been connected to data-driven and prescriptive decision–making. Smart predict-then-optimize and robust satisficing studies show that prediction errors should be evaluated through their downstream decision impact, not only through statistical loss [25, 26]. Related revenue management optimization work on static booking limit structure with buy-up behavior also illustrates the continuing role of optimization models in capacity control settings [42]. However, such models do not address hotel-specific cancellation calibration or aggregate residual load ambiguity in the sense considered here.

Our model assigns distinct roles to the two uncertainty sets. The calibration set perturbs the segment-specific location inputs that determine contractual value and mean load. The Wasserstein set perturbs the conditional residual load law and therefore its skewness, support, and tails. The robust CVaR constraint adds a third, managerial layer: It controls the tolerated severity of overflow rather than representing another ambiguity source. Keeping these layers separate makes the economic mechanism identifiable and allows the numerical analysis to ask whether the choice of uncertainty representation materially changes the resulting booking portfolio and cost profile.

2.3. *Research gap and positioning*

The reviewed literature leaves a precise optimization gap unresolved. Booking control models account for cancellations, no-shows, customer heterogeneity, risk aversion, and forecast uncertainty; prediction research provides increasingly refined reservation-level probabilities; and modern DRO protects decisions against ambiguity in probability laws. Yet these streams generally do not separate segment-level probability miscalibration from ambiguity in the conditional aggregate-load distribution. As a result, they do not characterize whether an adversarial probability perturbation should increase or decrease a segment's show-up probability once both contractual value and endogenous capacity scarcity are considered.

The gap addressed here is not merely that cancellation probabilities are heterogeneous. That fact is well recognized in hotel analytics. The unresolved optimization question is how a decision-maker should protect a booking portfolio when the heterogeneous probabilities themselves may be miscalibrated and when the economically adverse direction of miscalibration depends on the segment value, load contribution, and the endogenous marginal value of capacity. To the best of our knowledge, prior booking control models have not jointly represented segment-level calibration ambiguity and aggregate distributional ambiguity while deriving an endogenous reversal in the worst-case direction of probability error. The novelty is therefore not robust revenue management in general, nor Wasserstein DRO in isolation. It is the booking-specific structural link among calibration error, segment value-to-load ratios, and the marginal price of capacity. This link yields a local phase transition, an extreme-point adversary, a box-local calibration diagnostic, an active scenario weighted-budget certificate, and a measurable price of calibration uncertainty.

Table 1. Positioning of the proposed framework relative to related research streams.

Research stream	Decision variable	Treatment of probability error	Treatment of aggregate load ambiguity	Limitation relative to this paper
Classical and robust overbooking [3, 6, 39]	Scalar or class-level booking limit	Known rates, scenarios, or aggregate forecast error protection	Binomial, choice-based, or scenario-based demand	Strong capacity control structure, but limited separation between segment-level calibration error and aggregate load law ambiguity
Hotel revenue management and overbooking performance [3, 4, 33]	Room allocation, booking limits, or overbooking levels	Forecast accuracy and cancellation/no-show assumptions affect performance	Usually handled through demand models, historical rates, or managerial overbooking rules	Directly motivates the hotel setting, but does not derive bidirectional worst-case calibration effects
Cancellation prediction [7, 9–11, 38]	No explicit optimization decision, or heuristic downstream use	Reservation-level point estimates and calibration diagnostics	Usually not a central decision object	Motivates heterogeneous probabilities, but generally does not optimize a robust booking portfolio under miscalibration
Prediction-assisted booking control [12–14]	Scalar or segment-aware booking control	Estimated probabilities treated mainly as fixed inputs	Discrete, simulated, or forecast-based arrivals	Connects prediction to booking value, but does not characterize endogenous reversal of worst-case calibration direction
Revenue-management optimization and general robust/DRO models [15, 17, 23, 42]	Generic optimization decisions	Parameter uncertainty sets or distributional ambiguity sets	Application-dependent uncertainty or distributional model	Provides tractable protection tools, but does not identify the booking-specific value–load mechanism behind bidirectional calibration
Proposed framework	Multidimensional segment acceptance portfolio	Weighted budgeted calibration uncertainty with endogenous perturbation direction	Truncated skew-normal reference law protected by 1-Wasserstein ambiguity	Bidirectional calibration principle, segment phase transition, box-local marginal diagnostic, weighted-budget active scenario certificate, price of calibration uncertainty, finite convergence cutting-plane algorithm, and hotel cost-structured interpretation

Note: The comparison emphasizes the optimization decision, uncertainty representation, and booking-specific mechanism most directly related to the present study.

Table 1 summarizes the positioning. The proposed model differs from cancellation prediction because it does not take statistical accuracy as the final objective; it differs from classical overbooking because it protects a segment-level portfolio rather than only a scalar overbooking level; and it differs from general robust or distributionally robust optimization because the worst-case probability direction is characterized through a booking-specific value-to-load threshold. This positioning also clarifies the role of the newly added hotel and revenue management references: They strengthen the empirical and journal-specific motivation, while the methodological contribution remains the calibration-robust booking-control formulation and its structural analysis.

3. Model formulation

This section separates three conceptually distinct objects: Segment-level calibration ambiguity, ambiguity in the conditional residual load law, and managerial tail-risk tolerance. The first is modeled through a weighted budgeted set in the spirit of budgeted robust optimization [15]; the second through a

1-Wasserstein neighborhood supported by modern data-driven DRO theory [17, 18, 20]; and the third through a CVaR service constraint [27]. This separation is essential because the three parameters control different economic mechanisms rather than interchangeable degrees of conservatism.

Table 2 gives a compact dictionary for the symbols used in the formulation. It groups the notation by booking primitives, calibration uncertainty, residual load ambiguity, and numerical diagnostics so that later equations can be read without repeated definitions. The table is descriptive only; it does not introduce additional assumptions. The public hotel booking records used later provide only calibration range grounding for selected segment and cost proxies; they are not used as field evidence of implemented booking control performance.

3.1. Operational cost mechanism

The formulation is intended as an optimization model with a hotel cost-structured application, not as a new cancellation prediction model. Robustness enters the hotel decision problem through three operational channels. First, segment-level acceptance quantities determine which booking classes occupy scarce capacity. Second, calibration ambiguity protects the model against misspecified show-up rates within those accepted segments. Third, distributional ambiguity protects the aggregate residual load against tail misspecification after the segment portfolio has been chosen. Thus, the robust policy is not merely a uniformly more conservative overbooking rule. It reallocates the acceptance portfolio across segments by weighing contractual value, room-night load, vacancy cost, overflow cost, and service risk tolerance jointly.

In the hotel interpretation, the unused capacity penalty represents the opportunity cost of rooms that remain empty after the service date, whereas the overflow penalty represents walking, compensation, relocation, and reputational costs when realized arrivals exceed physical capacity. The robust CVaR constraint adds a service risk layer by limiting the tail-average severity of overflow rather than only its expected cost. These distinctions make the model directly interpretable as a cost-governed booking control problem.

Operationally, the mechanism is as follows. A calibration perturbation changes the plausible show-up rate of an accepted segment. This change affects both the expected contractual value and the expected room-night load of the segment. The optimizer then reallocates the accepted portfolio by comparing the segment's marginal value with its marginal capacity burden. This reallocation changes vacancy exposure, overflow exposure, and the robust service risk certificate. Robust optimization therefore affects hotel operations through segment selection and capacity protection, rather than through a uniform reduction in total accepted bookings.

3.2. Reservation portfolio and economic primitives

Consider one control epoch for a fixed service date or a fixed capacity cell. Reservations are partitioned into a finite set of segments

$$\mathcal{J} = \{1, \ldots, J\}. \tag{1}$$

A segment may represent a combination of property, room class, booking channel, lead time band, contract type, customer class, and stay pattern. For example, a segment can be written conceptually as $j = (k, c, h, t)$, where k denotes the room class, c denotes the booking channel, h denotes the lead-time band, and t denotes the stay load or contract class. The model does not require every physical room type, channel, or tariff rule to be listed separately in the formulation; these operational categories are encoded in the finite segment set $\mathcal{J}$ used by the booking control system.

Table 2. Main notation used in the calibration-robust booking control model.

Symbol	Type	Description
Panel A: Booking control primitives		
$j \in \mathcal{J}$	Index/set	Booking segment and segment set.
N_j	Parameter	Available booking requests in segment j.
x_j	Decision	Accepted bookings in segment j.
a_j	Parameter	Room-night load per accepted booking.
r_j	Parameter	Net contribution from a materialized arrival.
f_j	Parameter	Retained value from cancellation or nonrefundable terms.
C	Parameter	Physical capacity of the focal capacity cell.
c_o, c_u	Parameters	Unit overflow and unused capacity costs.
A, b	Parameters	Constraint matrix and right-hand side for additional linear operating restrictions.
$\mathcal{X}$	Set	Feasible integer acceptance portfolios.
Panel B: Show-up, calibration, and value load quantities		
$\widehat{q}_j$	Parameter	Nominal calibrated show-up rate.
q_j	Uncertain parameter	Perturbed show-up rate.
δ_j	Parameter	Segment-level calibration half-width.
u_j	Perturbation	Normalized displacement in $[-1, 1]$.
ω_j	Weight	Reliability or prevalence weight in the calibration budget.
Γ	Parameter	Weighted calibration budget.
$\mathcal{U}_{\mathrm{cal}}(\Gamma)$	Set	Budgeted calibration uncertainty set.
τ_j	Threshold	Value load threshold for direction reversal.
g	Marginal value	Shadow value in the direction reversal condition.
Panel C: Residual load ambiguity and risk certificate		
$m(x, q)$	Function	Expected service load.
Z	Random variable	Standardized residuaload disturbance.
$D(x, q, Z)$	Random variable	Realized aggregate service load.
$\sigma(x)$	Function	Portfolio-dependent residual load scale.
ℓ	Support bound	Lower support bound of the standardized residual law.
$\widehat{P}_{\alpha,\ell}$	Distribution	Fitted truncated skew-normal law for Z.
Q	Distribution	Adversarial residual load law.
ε	Parameter	Wasserstein ambiguity radius.
$\mathcal{P}_{\varepsilon}$	Set	Wasserstein ambiguity set around $\widehat{P}_{\alpha,\ell}$.
β, κ	Parameters	CVaR confidence level and permitted tail-average overflow.
CVaR_{β}	Risk measure	Service risk certificate.
Panel D: Robust values and diagnostics		
$R(x, q)$	Function	Expected contractual value.
$\mathcal{L}(d; C)$	Function	Asymmetric mismatch loss.
PCU	Metric	Price of calibration uncertainty.
$M_j^{\mathrm{box}}(g)$	Diagnostic	Box-local calibration diagnostic.
$Q^{\star}(x)$	Set	Active worst-case calibration scenarios.
Gap	Metric	Finite-model value gap.
Res	Metric	Feasibility residual.

For segment $j \in \mathcal{J}$, $N_j \in \mathbb{Z}_+$ denotes the available booking requests and $x_j \in \mathbb{Z}_+$ denotes the accepted quantity selected by the booking control policy. The variable x_j does not mean that hotel managers manually approve individual reservations after customers submit booking requests. It denotes the number of requests from segment j authorized for acceptance by the booking-control policy before that segment is closed, restricted, or protected for other demand. Operationally, this may be implemented through booking limits, room-class availability controls, rate-plan restrictions, channel allotments, or group exposure limits.

The parameter $a_j > 0$ maps one accepted request into capacity load, while $r_j > 0$ and $f_j \geq 0$ denote, respectively, the net value of a materialized arrival and the retained contractual value from a cancellation or nonrefundable booking. The realized, unknown show-up probability $q_j \in [0, 1]$ determines how the accepted portfolio translates into expected service load and is later allowed to vary within a calibration uncertainty set. These primitives are summarized in Table 2 and are used below to define the feasible acceptance set and the economic objective.

The feasible set is

$$\mathcal{X} = \left\{x \in \mathbb{Z}_+^J : 0 \leq x_j \leq N_j,\ j \in \mathcal{J},\ Ax \leq b\right\}. \tag{2}$$

The set in Equation (2) restricts the multidimensional acceptance portfolio through segment availability and any additional operational constraints represented by $Ax \leq b$. Here, A and b denote the matrix and right-hand side that define additional linear operating restrictions; they are not themselves separate constraints. The compact inequality $Ax \leq b$ may represent room-class nesting, channel allotments, group-block caps, contractual allotments, limits on high-load segments, or protection levels reserved for later arriving demand.

For a fixed probability vector q, the expected contractual revenue is

$$R(x, q) = \sum_{j \in \mathcal{J}} \left[q_j r_j + (1 - q_j) f_j\right] x_j. \tag{3}$$

Equation (3) separates the realized service value r_j from the retained cancellation value f_j. In a hotel cost application, a_j may represent room-night load, r_j is the net contribution of an arriving booking, and f_j is the retained value from a cancelled or nonrefundable booking. Thus, x_j is not a scalar overbooking buffer but a segment-level acceptance decision. The same show-up probabilities determine the expected service load, which is

$$m(x, q) = \sum_{j \in \mathcal{J}} a_j q_j x_j. \tag{4}$$

Definition 3.1 (Bidirectional calibration uncertainty)**.** Because Equations (3) and (4) depend on the same probability vector, a probability perturbation changes both the contractual value and capacity burden. Calibration uncertainty is bidirectional when the economically adverse perturbation of a segment's show-up probability may be either upward or downward, depending on the relative magnitudes of incremental service value and marginal capacity burden. The direction is determined endogenously by the optimization model rather than fixed ex ante.

3.3. Calibration uncertainty set

The preceding value load formulation explains why the direction of probability error cannot be fixed in advance. We therefore adapt the weighted budget construction of [15] to probability inputs while allowing the economically adverse sign to be determined endogenously by each segment's marginal value and capacity burden.

Let $\widehat{q}_j$ denote the nominal calibrated show-up rate for segment j. The true rate is not assumed known. Let $\delta_j \geq 0$ be a segment-specific calibration radius obtained from an out-of-time confidence interval, bootstrap procedure, conformal calibration band, or simultaneous binomial bound.

For later structural statements, define the feasible coordinate endpoints

$$\underline{q}_j = \max\{0, \widehat{q}_j - \delta_j\}, \qquad \overline{q}_j = \min\{1, \widehat{q}_j + \delta_j\}.$$

Thus the admissible coordinate interval is $I_j = [\underline{q}_j, \overline{q}_j]$. When the symmetric calibration band lies strictly inside $[0, 1]$, these endpoints reduce to $\widehat{q}_j \pm \delta_j$; otherwise, the probability bounds provide the required clipping.

We represent calibration error as

$$q_j = \widehat{q}_j + \delta_j u_j, \qquad -1 \leq u_j \leq 1. \tag{5}$$

Equation (5) separates the nominal rate from its admissible signed displacement. To avoid the excessive conservatism of simultaneous endpoint failure in all segments, define the weighted budgeted set

$$\mathcal{U}_{\mathrm{cal}}(\Gamma) = \left\{ q \in [0,1]^J : q_j = \widehat{q}_j + \delta_j u_j,\ |u_j| \leq 1,\ \sum_{j \in \mathcal{J}} \omega_j |u_j| \leq \Gamma \right\}, \tag{6}$$

where $\omega_j > 0$ is a reliability or prevalence weight and $\Gamma \geq 0$ is the calibration budget. Thus, Equation (6) limits joint miscalibration without fixing the adverse sign in advance.

Remark 3.2. The calibration set is not an evaluation score. It is an optimization object. The radius δ_j determines the admissible local displacement of segment j, whereas Γ controls how many or how strongly several segments may be jointly misspecified. In the hotel cost interpretation, Γ is a governance parameter for segment-level calibration risk rather than a generic conservatism parameter.

3.4. Standardized truncated skew-normal service load

Calibration ambiguity governs uncertainty in the segment-level probabilities that determine the mean service load. Even conditional on those probabilities, however, the realized aggregate load can deviate from its mean, which motivates a separate residual load model.

We use the skew-normal family as a flexible reference model for an asymmetric residual load and impose lower truncation to respect the operational support of realized service load; see [43] for the skew-normal family and its standardization properties. Let $\widetilde{Z}$ follow a truncated skew-normal distribution with the shape parameter α and the lower truncation point $\widetilde{\ell}$. Define the standardized disturbance

$$Z = \frac{\widetilde{Z} - \mathbb{E}[\widetilde{Z}]}{\sqrt{\mathrm{Var}(\widetilde{Z})}}. \tag{7}$$

By Equation (7), $\mathbb{E}[Z] = 0$, $\mathrm{Var}(Z) = 1$, and $\mathrm{supp}(Z) = [\ell, \infty)$ for an induced lower bound $\ell < 0$.

The realized aggregate service load is

$$D(x, q, Z) = m(x, q) + \sigma(x) Z, \tag{8}$$

where $\sigma(x) > 0$ is the residual scale. Equation (8) decomposes realized load into the probability-dependent mean in Equation (4) and a standardized residual component. The structural analysis first uses a fixed σ, while the computational extension allows

$$\sigma^2(x) = \sigma_0^2 + \sum_{j \in \mathcal{J}} v_j x_j + x^\top \Sigma x. \tag{9}$$

In the computational specification, we take $\sigma_0 > 0$, $v_j \geq 0$, and $\Sigma \succeq 0$, so that $\sigma(x)$ is well defined and strictly positive for every feasible portfolio. When the scale is portfolio dependent as in Equation (9), the admissible portfolio must also satisfy the support compatibility condition

$$m(x, q) + \sigma(x)\ell \geq 0 \qquad \forall q \in \mathcal{U}_{\mathrm{cal}}(\Gamma). \tag{10}$$

3.5. Asymmetric capacity loss

Given the realized load representation in Equation (8), the operational consequence of load uncertainty is determined by how deviations from physical capacity are valued.

Let $C > 0$ denote physical service capacity, $c_o > 0$ denote the unit cost of overflow or denied service, and $c_u > 0$ denote the unit cost of unused capacity. The asymmetric mismatch loss is

$$\mathcal{L}(d; C) = c_o \left(d - C\right)^+ + c_u \left(C - d\right)^+. \tag{11}$$

In Equation (11), the first term captures relocation, compensation, and service recovery, whereas the second captures spoilage of perishable capacity and forgone contribution.

3.6. Distributional ambiguity

Because the expected mismatch loss depends on the residual load law, protection against misspecification of that law forms the second uncertainty layer of the model. Wasserstein ambiguity is attractive because it preserves geometric information about the perturbations, admits finite-sample protection under suitable radius choices, and often yields tractable dual representations [17, 18, 20, 22]. The ambiguity radius ε therefore controls misspecification of the conditional residual load law rather than calibration error in the segment probabilities.

Let $\widehat{P}_{\alpha,\ell}$ be the fitted standardized truncated skew-normal law of Z. The residual law is modeled as lying in the 1-Wasserstein ball

$$\mathcal{P}_\varepsilon = \left\{ Q : W_1(Q, \widehat{P}_{\alpha,\ell}) \leq \varepsilon,\ \mathrm{supp}(Q) \subseteq [\ell, \infty) \right\}. \tag{12}$$

The radius $\varepsilon \geq 0$ in Equation (12) captures residual distributional misspecification separately from the calibration budget Γ in Equation (6). The two radii need not have the same operational effect: Γ changes the admissible segment-level show-up rates, whereas ε changes the protection against aggregate residual load tail behavior. This distinction motivates the uncertainty set sensitivity analysis in the computational section. The Wasserstein ball is placed on the standardized residual load disturbance Z. Through the affine load mapping $D(x, q, Z)$, this ambiguity induces protection against misspecification in the aggregate service load distribution for any fixed portfolio and calibration vector.

3.7. Calibration-aware distributionally robust model

Having specified segment-level calibration ambiguity and conditional residual load ambiguity separately, we now combine the two layers in a single booking control problem. This nested construction is consistent with the broader primal–dual unification of robust and distributionally robust optimization developed by [23].

The complete booking control formulation is

$$\begin{aligned} \max_{x\in\mathcal{X}} \quad & \inf_{q\in\mathcal{U}_{\mathrm{cal}}(\Gamma)} \left\{ R(x,q) - \sup_{Q\in\mathcal{P}_\varepsilon} \mathbb{E}_Q\left[\mathcal{L}(D(x,q,Z);C)\right] \right\} \\ & \text{s.t.} \quad \sup_{q\in\mathcal{U}_{\mathrm{cal}}(\Gamma)} \sup_{Q\in\mathcal{P}_\varepsilon} \mathrm{CVaR}_\beta^Q\left[(D(x,q,Z)-C)^+\right] \le \kappa, \\ & \qquad m(x,q) + \sigma(x)\ell \ge 0, \qquad \forall q\in\mathcal{U}_{\mathrm{cal}}(\Gamma). \end{aligned} \tag{13}$$

Optimization problem in Equation (13) combines the revenue expression in Equation (3), the mean load mapping in Equation (4), and the asymmetric loss in Equation (11) under two distinct uncertainty sets.

For later reformulation, it is useful to write the same model in value function form. Define

$$\Psi_{\Gamma,\varepsilon}(x) = \inf_{q\in\mathcal{U}_{\mathrm{cal}}(\Gamma)} \left\{ R(x,q) - \sup_{Q\in\mathcal{P}_\varepsilon} \mathbb{E}_Q[\mathcal{L}(D(x,q,Z);C)] \right\}, \tag{14}$$

$$\mathcal{X}_{\Gamma,\varepsilon} = \left\{ x\in\mathcal{X} : \sup_{q\in\mathcal{U}_{\mathrm{cal}}(\Gamma)} \sup_{Q\in\mathcal{P}_\varepsilon} \mathrm{CVaR}_\beta^Q[(D(x,q,Z)-C)^+] \le \kappa,\ \ m(x,q)+\sigma(x)\ell \ge 0\ \forall q\in\mathcal{U}_{\mathrm{cal}}(\Gamma) \right\}. \tag{15}$$

Then the optimization problem in Equation (13) is, equivalently

$$V(\Gamma,\varepsilon) = \max_{x\in\mathcal{X}_{\Gamma,\varepsilon}} \Psi_{\Gamma,\varepsilon}(x). \tag{16}$$

Equations (14)–(16) make explicit that the objective adversary and the service risk adversary are evaluated on the same accepted portfolio, but they enter the model through separate value and feasibility certificates. Here $\beta\in(0,1)$ is the tail confidence level and $\kappa\ge 0$ is the maximum permitted tail-average overflow.

The Optimization problem in Equatin (13) operationalizes robust optimization as a portfolio reallocation mechanism. A larger calibration budget may reduce acceptance of segments whose show-up errors create high capacity burden, but it may preserve or even favor segments whose value-to-load profile remains attractive under stress. Similarly, Wasserstein ambiguity does not directly alter the nominal segment rates; it protects the residual aggregate load law and can change how much capacity protection is needed for a given portfolio. This separation allows the numerical analysis to ask whether the uncertainty set specification changes not only the robust value but also the composition and cost profile of the optimal booking policy.

Remark 3.3. The optimization problem in Equation (13) has a max–min–sup structure with two distinct adversaries. The calibration adversary selects plausible segment-level show-up rates, while the distributional adversary selects a plausible residual law. Their roles should not be combined because they affect different economic primitives.

The formulation also explains why a separate structural analysis is needed. Because the same segment-level show-up probability affects both contractual value and expected capacity load, a calibration perturbation creates competing value and congestion effects. A downward perturbation may reduce contractual value while relieving capacity pressure, whereas an upward perturbation may increase value while intensifying overflow exposure. Section 4 characterizes this value load trade-off and identifies when the economically adverse calibration direction reverses.

This distinction is operationally important because a uniformly conservative adjustment can protect the wrong side of the uncertainty interval for some segments. When capacity is relatively slack, a lower show-up rate can be more damaging because it erodes contractual value without creating a material congestion benefit. As capacity becomes scarce, the same segment may instead be harmed by an upward show-up perturbation that intensifies overflow exposure. The relevant stress direction is therefore determined jointly by segment economics and the endogenous scarcity value of capacity. This mechanism motivates the threshold results developed next and provides the economic link between the booking-control formulation and the segment-specific calibration diagnostics used in the numerical study.

The same logic also clarifies why calibration robustness should be evaluated at the portfolio level rather than through isolated probability corrections. A perturbation that appears conservative for one segment can reallocate scarce capacity toward or away from other segments, thereby changing both the objective contribution and the binding service-risk margin. Consequently, the relevant robustness question is not simply whether each show-up probability is shifted upward or downward, but which joint perturbation is economically adverse for the accepted mix. The structural results below formalize this portfolio interaction through marginal capacity values, segment-specific reversal thresholds, and budget allocation across active calibration coordinates. These results then provide a direct interpretation of the active adversarial scenarios observed in the computational study.

4. Structural analysis

This section derives structural properties that are specific to calibration-aware booking control, not merely restatements of generic robust-optimization theory. Budgeted robustness explains why extreme calibration scenarios are sufficient [15], while optimal transport theory later supports the distributional counterpart. The central contribution is not the observation that cancellation risk is heterogeneous; that fact is already well recognized in hotel analytics. The contribution is to show that, under heterogeneous value load segments, the economically adverse calibration direction itself is endogenous: It is segment-dependent and can reverse as the marginal capacity price changes.

The results in this section have two roles. Computationally, they reduce the calibration adversary to extreme points, local diagnostics, and active scenarios used later in the finite reformulation and cutting-plane algorithm. Operationally, they explain when capacity protection should stress show-up probabilities upward or downward, how segment value-to-load ratios affect acceptance priorities, and why calibration robustness is segment selective rather than a uniform booking reduction.

Table 3 provides a compact dependency map for the main theoretical results and numerical diagnostics. The map is descriptive; it does not add assumptions beyond the primitives and uncertainty sets already defined in the model formulation. Its role is to help readers follow the theory-to-algorithm-to-diagnostic flow without repeating a separate synthetic design table in the computational section.

Table 3. Integrated theory-to-experiment dependency map.

Result group	Required primitives	Numerical diagnostic	Purpose in the manuscript
Nominal asymmetric loss; Theorem 4.1, Corollary 4.2, Proposition 4.3	Capacity C, costs c_o, c_u, standardized residual law, loss $\mathcal{L}(d; C)$	Reference law fit and capacity-loss geometry	Establishes the one-dimensional mismatch loss structure used to interpret target load, protection levels, and service load asymmetry.
Skewness, truncation, and convex-order effects; Proposition 4.4, Theorem 4.6	Truncated skew-normal reference law, asymmetric loss, lower-support restriction	Reference law fit, tail diagnostics, and Wasserstein-radius sensitivity	Motivates the asymmetric residual reference and then evaluates distributional protection through the ambiguity radius used in the robust model.
Bidirectional calibration and phase transition; Theorems 4.7–4.8	Segment value r_j, retained value f_j, load a_j, calibration radius δ_j, marginal capacity value g	Direction reversal and phase-transition diagnostics	Demonstrates that the adverse calibration direction is segment-dependent and can reverse endogenously as the marginal value of capacity changes.
Extreme-point adversary and budget allocation; Theorem 4.9, Proposition 4.10	Weighted calibration budget Γ, weights ω_j, segment perturbations u_j	Adversarial allocation diagnostic	Reduces calibration stress to interpretable active segment-level scenarios and identifies which segments consume the calibration budget.
Local diagnostic and active-scenario certificate; Proposition 4.11, Theorem 4.12	Box-local marginal diagnostic, active scenarios, weighted-budget Karush–Kuhn–Tucker (KKT) conditions	Local diagnostic and weighted budget certificate	Distinguishes coordinatewise segment diagnostics from the global weighted budget optimality certificate.
Price of calibration uncertainty and policy monotonicity; Theorems 4.13, 4.14, and 4.15	Calibration budget Γ, Wasserstein radius ε, robust objective, normalized policy primitives	Value of protection and uncertainty set sensitivity diagnostics	Interprets robustness as a measurable cost of calibration protection and distributional ambiguity, not as an arbitrary conservative adjustment.
Wasserstein reformulation and finite inner maximization; Theorem 5.1, Corollary 5.2, Propositions 5.3–5.4	Wasserstein radius ε, support $[\ell, \infty)$, Lipschitz capacity loss terms	Distributional-stress and robust-CVaR computations	Converts aggregate residual load ambiguity into a deterministic finite reformulation suitable for numerical certification.
Budgeted counterpart, cutting-plane convergence, and out-of-sample protection; Proposition 5.6, Theorems 5.8 and 5.9	Finite calibration scenarios, robust CVaR pieces, support constraints, statistical coverage events	Algorithm-efficiency and finite gap audit	Provides computational certification through iteration counts, active cuts, residuals, finite model gaps, and calibrated ambiguity set coverage.

4.1. Nominal asymmetric capacity loss

For fixed $\sigma > 0$ and a reference disturbance Z with continuous density f_Z, define

$$G(m) = \mathbb{E}\left[c_o (m + \sigma Z - C)^+ + c_u (C - m - \sigma Z)^+\right]. \tag{17}$$

Theorem 4.1 (Convexity and derivatives of asymmetric capacity loss)**.** *Suppose Z has a continuous distribution function F_Z and density f_Z. Then G is convex and continuously differentiable, with*

$$G'(m) = c_o - (c_o + c_u) F_Z\left(\frac{C - m}{\sigma}\right). \tag{18}$$

Wherever $f_Z((C - m)/\sigma) > 0$, G is twice differentiable and

$$G''(m) = \frac{c_o + c_u}{\sigma} f_Z\left(\frac{C - m}{\sigma}\right) > 0. \tag{19}$$

Equation (18) identifies the marginal capacity price under the reference law: Overflow probability raises the derivative through c_o, whereas unused capacity probability lowers it through c_u. Equation (19) then shows that the expected mismatch loss is strictly convex wherever the density is positive. Operationally, this marginal price is the capacity scarcity signal that later determines whether a segment's show-up probability is stressed upward or downward.

Corollary 4.2 (Unique nominal target load)**.** *If F_Z is strictly increasing on the relevant support, the unconstrained minimizer of G is unique and satisfies*

$$F_Z\left(\frac{C - m^\star}{\sigma}\right) = \frac{c_o}{c_o + c_u}. \tag{20}$$

Equivalently,

$$m^\star = C - \sigma F_Z^{-1}\left(\frac{c_o}{c_o + c_u}\right). \tag{21}$$

Setting the derivative in Equation (18) equal to zero yields the critical fractile condition in Equation (20); solving that condition for the mean load gives Equation (21). This benchmark describes the target load before calibration uncertainty is introduced, so later deviations can be interpreted as robustness-driven portfolio adjustments rather than changes in the basic cost trade-off.

Proposition 4.3 (Cost comparative statics)**.** *Under Corollary 4.2, $m^\star$ increases one-for-one with C. If F_Z^{-1} is increasing, then $m^\star$ weakly decreases with c_o and weakly increases with c_u.*

Proposition 4.3 gives the managerial comparative statics: Higher overflow cost makes the target load more conservative, whereas higher unused capacity cost makes the target load more aggressive.

4.2. Skewness and truncation

Let $F_{\alpha,\ell}$ denote the standardized truncated skew-normal distribution function. Define the critical quantile

$$z_u(\alpha, \ell) = F_{\alpha,\ell}^{-1}(u), \qquad u = \frac{c_o}{c_o + c_u}. \tag{22}$$

Proposition 4.4 (Local skewness effect)**.** *Suppose $F_{\alpha,\ell}$ is continuously differentiable in z and α, with positive density at z_u. Then*

$$\frac{\partial m^\star}{\partial \alpha} = \sigma \frac{\partial F_{\alpha,\ell}(z_u)/\partial \alpha}{f_{\alpha,\ell}(z_u)}. \tag{23}$$

Hence, Equation (23) *shows that the direction of the skewness effect is determined by the local distributional shift at the critical quantile in Equation* (22)*, rather than by the sign of α alone.*

Remark 4.5. The proposition avoids an unjustified global claim that greater positive skewness always lowers the optimal load. Standardization and truncation can change quantile ordering. The sign must be established analytically for a restricted parameter region or verified numerically at the relevant critical fractile.

4.3. Convex order and stop-loss dominance

The following comparison uses the standard implication that convex order preserves expectations of convex loss functions [44].

Theorem 4.6 (Capacity loss ordering)**.** *Let D_1 and D_2 have finite first moments and suppose $D_1 \leq_{\mathrm{cx}} D_2$. Then, for every capacity C, we have*

$$\mathbb{E}[(D_1 - C)^+] \leq \mathbb{E}[(D_2 - C)^+], \tag{24}$$

and

$$\mathbb{E}[\mathcal{L}(D_1; C)] \leq \mathbb{E}[\mathcal{L}(D_2; C)]. \tag{25}$$

Equation (24) compares overflow exposure directly, while Equation (25) extends the ordering to the full asymmetric mismatch loss. Theorem 4.6 therefore shows that mean service load alone is insufficient for booking control: A more dispersed distribution generates larger expected loss even when its mean is unchanged.

4.4. Bidirectional calibration and direction reversal

For a fixed x, and hence a fixed residual scale $\sigma = \sigma(x)$, define the one-dimensional robust capacity loss function

$$H_\varepsilon(m) = \sup_{Q\in\mathcal{P}_\varepsilon} \mathbb{E}_Q[\mathcal{L}(m+\sigma Z; C)]\,. \tag{26}$$

Because $m \mapsto \mathcal{L}(m+\sigma z; C)$ is convex for every fixed z, H_ε is convex as the supremum of convex expected loss functions. For fixed x and q, define

$$\Phi(x,q) = R(x,q) - H_\varepsilon(m(x,q)). \tag{27}$$

Because Equation (27) combines contractual value and capacity loss through the common load mapping in Equation (4), the sign of a probability perturbation depends on both effects. Define the segment value-to-load threshold

$$\tau_j = \frac{r_j - f_j}{a_j}. \tag{28}$$

Theorem 4.7 (Local calibration-direction reversal)**.** *Suppose H_ε is differentiable at $m = m(x,q)$ and set $g = H_\varepsilon'(m(x,q))$. The directional derivative of the robust objective with respect to the coordinate q_j is*

$$\frac{\partial\Phi(x,q)}{\partial q_j} = x_j[(r_j - f_j) - a_j g]. \tag{29}$$

Equivalently, for any admissible signed perturbation t in coordinate j,

$$\Phi(x, q+te_j) - \Phi(x,q) = t x_j[(r_j - f_j) - a_j g] + o(t), \qquad t \to 0, \tag{30}$$

where e_j is the jth unit vector. Thus the sign of the coefficient in (30) *determines whether the adversary locally lowers or raises q_j. The locally adverse direction points toward*

$$q_j^{\text{local}}(g) = \begin{cases} \underline{q}_j, & \tau_j > g, \\ \overline{q}_j, & \tau_j < g, \end{cases} \tag{31}$$

with first-order neutrality when $\tau_j = g$. Moreover, if H_ε is differentiable throughout the load interval induced by $q_j \in I_j = [\underline{q}_j, \overline{q}_j]$ and $g(m) = H_\varepsilon'(m)$, then the lower endpoint is a global coordinatewise minimizer if $(r_j - f_j) - a_j g(m(x,q)) \geq 0$ throughout I_j, whereas the upper endpoint is a global coordinatewise minimizer if this quantity is nonpositive throughout I_j. If the sign changes inside I_j, the two endpoints must be compared directly.

Equation (29) isolates the two competing effects of q_j: $(r_j - f_j)$ is the incremental contractual value of realization, whereas $a_j g$ is the marginal capacity burden. The threshold $g = \tau_j$ therefore identifies a local reversal in the adverse perturbation direction. It gives an exact endpoint rule whenever the sign remains unchanged over the entire calibration interval, including any affine piece of a piecewise-linear loss representation. At a nondifferentiable load point, the later active-scenario KKT certificate uses the full subdifferential; Theorem 4.7 does not identify an arbitrary subgradient with the coordinate derivative.

This result clarifies the marginal contribution of heterogeneous cancellation risk in the present model. Heterogeneity matters not simply because different segments have different show-up rates, but because the same probability error can have opposite economic meanings across segments. A high-value, low-load segment and a low-value, high-load segment may be stressed in opposite directions under the same capacity state. In hotel terms, robust optimization may protect capacity by reducing exposure to segments whose show-up errors create high overflow cost, while retaining segments whose value-to-load ratios remain favorable even under calibration stress. This is the structural reason the robust policy cannot, in general, be replicated by a uniform cancellation rate stress test.

Theorem 4.8 (Local segment phase transition)**.** *Let $g(\theta)$ be a continuous marginal capacity price indexed by an operating parameter θ, such as the load pressure or the overflow-to-unused capacity cost ratio. The local adverse direction for segment j changes sign only when*

$$g(\theta) = \tau_j. \tag{32}$$

If the sign of $(r_j - f_j) - a_j g(\theta)$ is uniform over the relevant calibration interval on each side of the crossing, the corresponding global coordinatewise adverse endpoint also switches across this boundary.

The phase-transition statement gives a reviewer-facing interpretation of the mechanism. Capacity scarcity, service–cost asymmetry, and segment value jointly determine whether the adversary lowers or raises a segment's show-up probability. Thus the model generates a structural acceptance logic rather than a black-box conservative adjustment.

4.5. Extreme-point adversary and budget allocation

Theorem 4.9 (Extreme-point calibration adversary)**.** *For a fixed x, suppose H_ε is convex. Then $q \mapsto \Phi(x, q)$ is concave on $\mathcal{U}_{\mathrm{cal}}(\Gamma)$. If $\mathcal{U}_{\mathrm{cal}}(\Gamma)$ is a nonempty compact polytope, then*

$$\min_{q \in \mathcal{U}_{\mathrm{cal}}(\Gamma)} \Phi(x, q) = \min_{q \in \mathrm{ext}(\mathcal{U}_{\mathrm{cal}}(\Gamma))} \Phi(x, q). \tag{33}$$

For a local marginal price g, define the adversarial sensitivity score

$$s_j(x, g) = \delta_j x_j \left|(r_j - f_j) - a_j g\right|. \tag{34}$$

Proposition 4.10 (Adversarial budget allocation)**.** *In the linearized inner calibration problem, the adversary allocates the weighted budget to segments in nonincreasing order of*

$$\frac{s_j(x, g)}{\omega_j}, \tag{35}$$

up to each segment's feasible directional cap induced by $q_j \in [0, 1]$. There is an optimal allocation with at most one segment partially stressed relative to its feasible cap, apart from equivalent redistributions among tied scores.

Equation (34) shows that calibration exposure depends jointly on accepted quantity, calibration radius, and economic sensitivity. The weighted ranking in Equation (35) therefore implies that a large radius alone does not identify the most consequential segment. Operationally, the score identifies which accepted segments consume the common calibration budget in the worst case; this provides a direct bridge between the mathematical adversary and the hotel manager's exposure to segment-specific miscalibration.

4.6. Local box diagnostic and weighted-budget optimality certificate

For independent box calibration uncertainty, define the coordinatewise worst-case marginal contribution

$$M_j^{\text{box}}(g) = \min_{q_j \in [\widehat{q}_j - \delta_j, \widehat{q}_j + \delta_j] \cap [0,1]} \left\{ f_j + q_j[(r_j - f_j) - a_j g] \right\}. \tag{36}$$

Ignoring probability clipping only for notational clarity, we have

$$M_j^{\text{box}}(g) = f_j + \widehat{q}_j[(r_j - f_j) - a_j g] - \delta_j \left|(r_j - f_j) - a_j g\right|. \tag{37}$$

Proposition 4.11 (Box-local marginal diagnostic)**.** *Consider the continuous relaxation under independent box calibration uncertainty and suppose the active worst-case scenario is locally unique. At a candidate aggregate load with a marginal price g, segment j is locally attractive if $M_j^{\text{box}}(g) > 0$, locally unattractive if $M_j^{\text{box}}(g) < 0$, and locally neutral at equality.*

The result in Proposition 4.11 is intentionally local and box-specific. It is useful for interpreting the direction and magnitude of coordinatewise calibration exposure, but it is not a global acceptance rule for the weighted budgeted set $\mathcal{U}_{\text{cal}}(\Gamma)$. Under weighted budgeted ambiguity, the adversary allocates a common budget across coordinates, and several active worst-case extreme points may jointly support the robust optimum.

For the structural optimality certificate in this subsection, the residual scale is treated as fixed with respect to the acceptance vector, consistently with the fixed-scale structural specification used above. If the computational extension $\sigma = \sigma(x)$ is used instead, the stationarity condition must additionally include the derivative or subgradient contribution generated by the scale channel.

Let

$$Q^\star(x) = \operatorname*{arg\,min}_{q \in \text{ext}(\mathcal{U}_{\text{cal}}(\Gamma))} \Phi(x, q) \tag{38}$$

be the active worst-case calibration scenarios at x. For each $q \in Q^\star(x)$, define

$$G_j(x, q) = f_j + q_j[(r_j - f_j) - a_j g(x, q)], \qquad g(x, q) \in \partial H_\varepsilon(m(x, q)). \tag{39}$$

Theorem 4.12 (Weighted-budget active-scenario reduced KKT certificate)**.** *Consider the fixed-scale continuous relaxation with the bounds $0 \le x_j \le N_j$ and no active coupling, robust CVaR, or support compatibility constraint at $x^\star$. Suppose the robust objective is concave and the usual constraint qualification holds. At any robust optimum $x^\star$, there are non-negative weights $\{\lambda_q : q \in Q^\star(x^\star)\}$ summing to one such that*

$$\bar{G}_j = \sum_{q \in Q^\star(x^\star)} \lambda_q G_j(x^\star, q) \tag{40}$$

satisfies

$$\bar{G}_j \le 0 \quad \text{if } x_j^\star = 0, \qquad \bar{G}_j = 0 \quad \text{if } 0 < x_j^\star < N_j, \qquad \bar{G}_j \ge 0 \quad \text{if } x_j^\star = N_j. \tag{41}$$

If the active scenario is unique, the certificate reduces to the scenario marginal $G_j(x^\star, q^\star)$. When a coupling constraint, the robust CVaR constraint, or the support compatibility constraint is active, its multiplier-weighted subgradient must be added to the stationarity condition. The box-local expression in Equation (37) *is recovered only in the special coordinatewise box case.*

Minimizing the affine expression in Equation (36) over a single interval gives the absolute-value discount in Equation (37). Theorem 4.12, by contrast, gives the correct global first-order certificate for weighted budgeted ambiguity through the convex hull of active worst-case scenario gradients, using Danskin-type value-function calculus and normal-cone optimality [45, 46].

The certificate also improves auditability. It explains which active adverse scenarios support the robust booking decision and how their weighted marginal contributions justify zero, interior, or upper-bound acceptance of each segment. The robust policy is therefore interpretable as a certified segment portfolio rather than as an opaque reduction of total bookings.

4.7. Price of calibration uncertainty

Let $V(\Gamma, \varepsilon)$ denote the optimal value of the optimization problem (13). Define the price of calibration uncertainty (PCU) as

$$\mathrm{PCU}(\Gamma; \varepsilon) = V(0, \varepsilon) - V(\Gamma, \varepsilon). \tag{42}$$

Theorem 4.13 (Price of calibration uncertainty)**.** *For any fixed ε and any* $0 \le \Gamma_1 \le \Gamma_2$ *for which the corresponding robust problems are feasible, we have*

$$V(\Gamma_2, \varepsilon) \le V(\Gamma_1, \varepsilon).$$

Hence $\mathrm{PCU}(\Gamma; \varepsilon)$ *is non-negative and nondecreasing in* Γ.

For a fixed portfolio x, define the objective-only calibration exposure

$$\mathrm{PCU}_x(\Gamma; \varepsilon) = \Psi_{0,\varepsilon}(x) - \Psi_{\Gamma,\varepsilon}(x).$$

If H_ε is L_H-Lipschitz, then

$$0 \le \mathrm{PCU}_x(\Gamma; \varepsilon) \le \Gamma \max_{j \in \mathcal{J}} \frac{\delta_j x_j (|r_j - f_j| + a_j L_H)}{\omega_j} \le \Gamma \max_{j \in \mathcal{J}} \frac{\delta_j N_j (|r_j - f_j| + a_j L_H)}{\omega_j}. \tag{43}$$

The bound in Equation (43) *is a fixed-portfolio sensitivity bound. It does not, by itself, upper-bound the global PCU when increasing* Γ *also changes the robust feasible set. For the asymmetric capacity loss in Equation* (11)*, one may take* $L_H \le \max\{c_o, c_u\}$ *under location shifts.*

Thus PCU is the optimization-based price of guarding against segment-level calibration error. In the hotel application, it measures the value sacrificed to protect against show-up miscalibration while holding the residual load ambiguity radius fixed. This interpretation is useful for managerial reporting because it separates the cost of calibration protection from the cost of distributional protection.

4.8. Normalized policy equivalence

Define normalized primitives

$$\bar{a}_j = \frac{a_j}{C}, \qquad \bar{r}_j = \frac{r_j}{c_u C}, \qquad \bar{f}_j = \frac{f_j}{c_u C}, \qquad \rho = \frac{c_o}{c_u}, \qquad \bar{\sigma} = \frac{\sigma}{C}. \tag{44}$$

Theorem 4.14 (Normalized policy equivalence)**.** *Under the fixed-scale structural specification, consider two reservation capacity systems with identical feasible sets after capacity normalization and identical values of*

$$\left\{\bar{a}_j, \bar{r}_j, \bar{f}_j, \widehat{q}_j, \delta_j, \omega_j, \Gamma, \rho, \bar{\sigma}, \alpha, \ell, \varepsilon, \beta, \kappa/C\right\}. \tag{45}$$

Then their continuous relaxations have identical normalized optimal portfolios and identical normalized robust objective values. Thus the booking logic is invariant to a common scaling of capacity and monetary units. Because the Wasserstein ball in Equation (12) *is defined on the standardized residual Z, its radius ε is already expressed in standardized residual units and is therefore not divided by capacity. If the portfolio-dependent extension $\sigma = \sigma(x)$ is used, the same equivalence requires the entire normalized scale map $x \mapsto \sigma(x)/C$ to coincide across the two systems.*

4.9. Monotonicity of robust protection

Set inclusion immediately implies the monotonicity of the robust value. Continuity and optimizer existence results for Wasserstein model risk provide complementary regularity as the ambiguity radius varies [24].

Theorem 4.15 (Joint uncertainty monotonicity)**.** *If $\Gamma_2 \geq \Gamma_1$ and $\varepsilon_2 \geq \varepsilon_1$, then*

$$V(\Gamma_2, \varepsilon_2) \leq V(\Gamma_1, \varepsilon_1). \tag{46}$$

Remark 4.16. The monotonicity result concerns the robust objective value. The monotonicity of every component of the optimal acceptance vector does not follow without additional single-crossing or lattice conditions because uncertainty may change the composition of the optimal portfolio. Consequently, uncertainty set sensitivity should be evaluated not only through objective values but also through accepted load, overflow exposure, vacancy cost, and segment-level portfolio shifts.

4.10. Calibration sensitivity and decision stability

For a fixed portfolio x and residual law Q, define

$$\Pi(x; q, Q) = R(x, q) - \mathbb{E}_Q[\mathcal{L}(D(x, q, Z); C)]. \tag{47}$$

Also define the fixed-portfolio expected capacity loss map

$$\mathcal{G}_Q(m) = \mathbb{E}_Q[\mathcal{L}(m + \sigma(x)Z; C)],$$

and let $L_{\mathcal{G}}$ be a Lipschitz constant of $\mathcal{G}_Q$ with respect to m. For the asymmetric capacity loss in Equation (11), one may take $L_{\mathcal{G}} \leq \max\{c_o, c_u\}$.

Proposition 4.17 (Policy value sensitivity)**.** *For any $q, q' \in [0, 1]^J$,*

$$|\Pi(x; q, Q) - \Pi(x; q', Q)| \leq L_q(x)\|q - q'\|_1, \tag{48}$$

where one admissible constant is

$$L_q(x) = \max_{j \in \mathcal{J}} \left\{x_j|r_j - f_j| + a_j x_j L_{\mathcal{G}}\right\}. \tag{49}$$

This result is used only to establish robustness of the optimization value. It is not developed as a forecast regret decomposition or retained value analysis. Its role is to show that small calibration changes have controlled value effects for a fixed portfolio, while the preceding results explain why optimal policies may still change composition when uncertainty sets or marginal capacity prices change.

5. Distributionally robust reformulation and solution method

This section converts the nested max–min–sup formulation into finite optimization components. Data-driven Wasserstein DRO supplies the empirical ambiguity set foundation [17]. General Wasserstein duality supports the worst-case expectation reformulation [20, 22], while the unified primal-worst–dual-best principle clarifies why finite convex reductions are available [23]. The service risk constraint uses the convex CVaR representation of [27], and the calibration counterpart uses the polyhedral budget structure of [15].

The structural results in Section 4 enter the computation in two direct ways. The extreme-point characterization guarantees that an adverse calibration vector can be taken at an extreme point of the budgeted set, while the weighted budget structure permits that extreme point to be generated by an exact linear separation oracle rather than by materializing the entire extreme-point set. The piecewise-linear capacity loss and CVaR representations then reduce the distributional layer to finitely many affine pieces. These reductions provide the computational role of the structural propositions; their operational role remains the interpretation of capacity scarcity, segment priority, and direction reversal.

5.1. Wasserstein dual representation

The following representation is a specialization of modern Wasserstein DRO duality to the one-dimensional conditional load disturbance. General strong duality results allow measurable or upper-semicontinuous losses under suitable growth and moment conditions [20, 22].

For a fixed (x, q), define

$$h_{x,q}(z) = \mathcal{L}(m(x, q) + \sigma(x)z; C). \tag{50}$$

The worst-case expected capacity loss is

$$\mathcal{G}_\varepsilon(x, q) = \sup_{Q\in\mathcal{P}_\varepsilon} \mathbb{E}_Q[h_{x,q}(Z)]. \tag{51}$$

Theorem 5.1 (Wasserstein dual reformulation)**.** *Under standard upper-semicontinuity and growth conditions, we have*

$$\mathcal{G}_\varepsilon(x, q) = \inf_{\lambda\geq 0} \left\{ \lambda\varepsilon + \mathbb{E}_{\widehat{P}_{\alpha,\ell}} \left[\sup_{z\geq\ell} \left\{ h_{x,q}(z) - \lambda|z - Z| \right\} \right] \right\}. \tag{52}$$

Equation (52) replaces the infinite-dimensional search over the probability laws in Equation (51) with a scalar transport multiplier and a pointwise adversarial maximization. This is the key step that makes the Wasserstein component computationally accessible.

5.2. Lipschitz bound

The function $h_{x,q}$ is piecewise linear in z with slopes $-c_u\sigma(x)$ and $c_o\sigma(x)$. Hence

$$\mathrm{Lip}(h_{x,q}) = \sigma(x) \max\{c_o, c_u\}. \tag{53}$$

Corollary 5.2 (Safe robust upper bound)**.** *For every* (x, q)*,*

$$\mathcal{G}_\varepsilon(x, q) \leq \mathbb{E}_{\widehat{P}_{\alpha,\ell}}[h_{x,q}(Z)] + \varepsilon\sigma(x) \max\{c_o, c_u\}. \tag{54}$$

The slope calculation in Equation (53) inserted into the Wasserstein perturbation bound yields Equation (54). The resulting additive term is an interpretable distributional risk charge: It increases with both the ambiguity radius and the sensitivity of the capacity loss to transport perturbations, consistent with the variation regularization interpretation of Wasserstein DRO [21].

Proposition 5.3 (Fixed-scale exact transport charges)**.** *Suppose that* $\sigma(x) \equiv \sigma > 0$*,* $c_o \geq c_u$*, the reference law* $\widehat{P}_{\alpha,\ell}$ *has a finite first moment, the Wasserstein ball is nonempty, and the common support* $[\ell, \infty)$ *is unbounded above. Then the upper-tail slope* σc_o *is reachable and*

$$H_\varepsilon(m) = H_0(m) + \varepsilon\sigma c_o. \tag{55}$$

Under the same fixed-scale specialization, and provided that the Rockafellar–Uryasev representation admits minimax interchange under the stated integrability and support conditions [20, 22, 27],

$$\sup_{Q\in\mathcal{P}_\varepsilon} \mathrm{CVaR}_\beta^Q[(m+\sigma Z - C)^+] = \mathrm{CVaR}_\beta^{\widehat{P}_{\alpha,\ell}}[(m+\sigma Z - C)^+] + \frac{\varepsilon\sigma}{1-\beta}.wehave \tag{56}$$

The numerical experiments use Proposition 5.3. The reference law terms are evaluated by deterministic finite-quantile quadrature and convex piecewise-linear interpolation; consequently, the cutting-plane certificate is exact for that explicitly discretized model, not for an undiscretized continuous distribution.

5.3. Finite-candidate inner maximization

The kink of $h_{x,q}$ occurs at

$$z_C(x,q) = \frac{C - m(x,q)}{\sigma(x)}. \tag{57}$$

For a sample realization $\widehat{z}_s$, the dual inner function

$$\psi_s(z; x, q, \lambda) = h_{x,q}(z) - \lambda|z - \widehat{z}_s| \tag{58}$$

is piecewise affine with breakpoints at ℓ, $\widehat{z}_s$, and the capacity kink $z_C(x,q)$ defined in equation (57). Under the bounded-slope condition $\lambda \geq \sigma(x)c_o$, the supremum is finite and can be obtained from these breakpoints.

Proposition 5.4 (Finite-candidate evaluation)**.** *For* $\lambda \geq \sigma(x)\max\{c_o, c_u\}$*, the inner supremum in* (52) *is attained at a point in*

$$\{\ell, \widehat{z}_s, \max\{\ell, z_C(x,q)\}\}. \tag{59}$$

Proposition 5.4 eliminates a continuous inner search: For each reference realization, the Wasserstein separation step needs only the support boundary, the reference point, and the capacity kink.

5.4. Robust CVaR reformulation

CVaR is used because it is coherent, convex, and admits an auxiliary variable representation suitable for robust reformulation [27]. Related Wasserstein risk constraint research illustrates the associated feasibility–optimality trade-off and out-of-sample interpretation [47].

Let

$$O(x,q,Z) = (D(x,q,Z) - C)^+. \tag{60}$$

Applying the Rockafellar–Uryasev representation to the overflow variable in Equation (60) gives

$$\mathrm{CVaR}^{Q}_{\beta}(O) = \min_{\tau\ge 0}\left\{\tau + \frac{1}{1-\beta}\mathbb{E}_{Q}[(O-\tau)^{+}]\right\}. \tag{61}$$

The positive-part term in Equation (61) is also piecewise linear and Lipschitz in Z. It therefore admits the same Wasserstein dual treatment as Equation (52).

Theorem 5.5 (Robust CVaR counterpart)**.** *Suppose that, for every fixed $q \in \mathcal{U}_{\mathrm{cal}}(\Gamma)$, the Rockafellar–Uryasev auxiliary minimization and the Wasserstein supremum admit the required minimax interchange. Then the robust CVaR constraint in the optimization problem* (13) *is equivalent to requiring that for every $q \in \mathcal{U}_{\mathrm{cal}}(\Gamma)$, a scenario-specific threshold $\tau_q \ge 0$ exists such that*

$$\tau_q + \frac{1}{1-\beta}\sup_{Q\in\mathcal{P}_{\varepsilon}}\mathbb{E}_{Q}[\big(O(x,q,Z)-\tau_q\big)^{+}] \le \kappa. \tag{62}$$

For each fixed q, the distributional supremum admits a finite-dimensional Wasserstein dual reformulation analogous to Theorem 5.1. A single common threshold across all calibration scenarios is not imposed; doing so would, in general, produce a stronger condition than the original robust CVaR constraint.

5.5. Budgeted calibration counterpart

The counterpart below follows the linear program duality for weighted budgeted uncertainty [15]. Its booking-specific feature is that the coefficient magnitude in Equation (63) is generated by segment economics and therefore changes with the affine loss or service-risk piece being separated.

For a fixed linear coefficient vector $\varphi \in \mathbb{R}^{J}$, consider first the interior-band case $[\widehat{q}_j - \delta_j, \widehat{q}_j + \delta_j] \subseteq [0,1]$ for every j. The budgeted worst case

$$\min_{q\in\mathcal{U}_{\mathrm{cal}}(\Gamma)} \varphi^{\top} q \tag{63}$$

can be represented by a linear program in the perturbations u_j. Its dual introduces one budget price and segment-specific deviation prices.

Proposition 5.6 (Linear robust counterpart for calibration uncertainty)**.** *Let $q_j = \widehat{q}_j + \delta_j u_j$, $|u_j| \le 1$, and $\sum_j \omega_j |u_j| \le \Gamma$, with the symmetric calibration bands contained in* $[0,1]$. *Then*

$$\min_{q\in\mathcal{U}_{\mathrm{cal}}(\Gamma)} \varphi^{\top} q = \varphi^{\top}\widehat{q} - \min_{\pi\ge 0,\ s_j\ge 0}\left\{\Gamma\pi + \sum_{j\in\mathcal{J}} s_j : \omega_j\pi + s_j \ge \delta_j|\varphi_j|\right\}. \tag{64}$$

Equation (64) replaces the calibration adversary by a budget price π and segment-specific deviation prices s_j. If a calibration band is clipped by the probability boundary, the same linear programming argument applies after replacing the symmetric unit deviation cap by the corresponding feasible directional cap. The numerical instances used below satisfy the interior band condition, so no probability clipping is active in the reported computations.

5.6. Exact scenario-and-piece cutting-plane solution

The finite computational model uses the fixed-scale specialization of Proposition 5.3, deterministic quantile quadrature, and finite piecewise linear interpolation for the reference capacity loss and overflow CVaR functions. Under this representation, the continuous relaxation is a linear program (LP); the same counterpart becomes a mixed-integer linear program (MILP) if integer booking limits are imposed. The reported numerical experiments use the continuous relaxation, while the finite reformulation itself is valid for either choice of booking variable integrality.

Let

$$\mathcal{E} = \operatorname{ext}(\mathcal{U}_{\mathrm{cal}}(\Gamma))$$

denote the finite extreme-point set of the calibration polytope, and let $\mathcal{P}_L$ and $\mathcal{P}_C$ denote the finite affine-piece sets for the reference mismatch loss and overflow CVaR functions. Write a loss piece as $h_p(m) = \alpha_p m + b_p$ and a CVaR piece as $c_h(m) = \gamma_h m + d_h$. Conceptually, the full finite master is

$$\max_{x\in\mathcal{X},\,\theta} \quad \theta \tag{65}$$

$$\text{s.t.} \quad \theta \le R(x,q^e) - \alpha_p m(x,q^e) - b_p - \varepsilon\sigma c_o, \qquad e\in\mathcal{E},\ p\in\mathcal{P}_L, \tag{66}$$

$$\gamma_h m(x,q^e) + d_h + \frac{\varepsilon\sigma}{1-\beta} \le \kappa, \qquad e\in\mathcal{E},\ h\in\mathcal{P}_C, \tag{67}$$

$$m(x,q^e) + \sigma\ell \ge 0, \qquad e\in\mathcal{E}. \tag{68}$$

The implementation does not materialize $\mathcal{E}$. For each affine piece, the calibration part of the separation problem is linear in q and is solved exactly by the weighted budget oracle implied by Proposition 5.6; an optimal oracle solution can be chosen at an extreme point by Theorem 4.9. Thus the algorithm generates only the adverse scenario–piece combinations that are needed by the incumbent solution.

For a restricted-master solution (x^K,θ^K), define the finite-model robust value

$$\underline{V}(x^K) = \min_{e\in\mathcal{E},\,p\in\mathcal{P}_L} \left\{R(x^K,q^e) - \alpha_p m(x^K,q^e) - b_p - \varepsilon\sigma c_o\right\}, \tag{69}$$

the largest robust CVaR violation

$$\Delta^K_{\mathrm{cvar}} = \left[\max_{e\in\mathcal{E},\,h\in\mathcal{P}_C} \left\{\gamma_h m(x^K,q^e) + d_h + \frac{\varepsilon\sigma}{1-\beta}\right\} - \kappa\right]_+, \tag{70}$$

and the support violation

$$\Delta^K_{\mathrm{sup}} = \left[-\min_{e\in\mathcal{E}} \left\{m(x^K,q^e) + \sigma\ell\right\}\right]_+. \tag{71}$$

The robust-value violation is

$$\Delta^K_{\mathrm{val}} = \left[\theta^K - \underline{V}(x^K)\right]_+. \tag{72}$$

Each positive residual identifies a precise omitted full-model constraint rather than a heuristic stress scenario.

Proposition 5.7 (Exact finite-model separation certificate)**.** *For any restricted-master solution (x^K,θ^K), the three residuals Δ^K_{val}, Δ^K_{cvar}, and Δ^K_{sup} are zero if and only if (x^K,θ^K) satisfies every value cut, every robust-CVaR cut, and every support constraint in the full finite master* (65)–(68)*. If all three residuals are at most ϵ_{alg}, then the incumbent is feasible for the full finite robust system up to the residual ϵ_{alg} and θ^K overestimates the finite robust value at x^K by at most ϵ_{alg}.*

Algorithm 1 Certified scenario-and-piece cutting-plane algorithm for the finite model.

Require: Calibration budget data; finite piece sets $\mathcal{P}_L, \mathcal{P}_C$; tolerance $\epsilon_{\text{alg}} > 0$
Ensure: Robust acceptance portfolio $x^\star$ and certificate $(\Delta_{\text{val}}, \Delta_{\text{cvar}}, \Delta_{\text{sup}})$

Initialize the restricted master with the nominal calibration scenario and its active loss, CVaR, and support constraints
repeat
  Solve the current LP/MILP restricted master to global optimality; obtain (x^K, θ^K)
  For each loss piece, solve the weighted budget separation problem exactly and retain the smallest value in (69)
  For each CVaR piece, solve the weighted budget separation problem exactly and retain the largest value in (70)
  Solve the weighted budget support separation problem in (71)
  Compute Δ^K_{val}, Δ^K_{cvar}, and Δ^K_{sup} and add every newly violated scenario–piece or support constraint
until $\max\{\Delta^K_{\text{val}}, \Delta^K_{\text{cvar}}, \Delta^K_{\text{sup}}\} \le \epsilon_{\text{alg}}$
return $x^\star \leftarrow x^K$ with the three-residual finite-model certificate

Theorem 5.8 (Finite convergence of the finite piecewise-linear model)**.** *Suppose that $\mathcal{U}_{\text{cal}}(\Gamma)$ has finitely many extreme points, the reference loss and CVaR approximations have finitely many affine pieces, and every LP or MILP master is solved exactly. Then Algorithm 1 terminates after finitely many iterations at a globally optimal solution of the full finite piecewise-linear robust model. The total number of distinct value, CVaR, and support constraints is at most*

$$|\mathcal{E}|(|\mathcal{P}_L| + |\mathcal{P}_C| + 1).$$

The theorem concerns the explicitly discretized numerical model. Accordingly, the numerical study reports the master status, generated cuts, robust value gaps, CVaR violations, and support violations directly. The matched benchmark in Section 7.7 additionally compares delayed generation with the monolithic finite robust dual counterpart under the same optimization backend, thereby separating formulation-level computational savings from solver brand effects.

5.7. *Out-of-sample protection*

Data-driven Wasserstein DRO can provide finite-sample coverage and out-of-sample performance, which guarantees when the ambiguity radius is selected from concentration bounds or resampling procedures [17, 18]. The theorem below combines such distributional coverage with coverage of the calibration set through a union bound.

Let q^0 and P^0 denote the unknown true show-up vector and residual distribution.

Theorem 5.9 (Joint out-of-sample protection)**.** *Suppose*

$$\mathbb{P}\left(q^0 \in \mathcal{U}_{\text{cal}}(\Gamma)\right) \ge 1 - \eta \tag{73}$$

and

$$\mathbb{P}\left(P^0 \in \mathcal{P}_\varepsilon\right) \ge 1 - \zeta. \tag{74}$$

Then, with a probability of at least $1-\eta-\zeta$, the simultaneous coverage events in Equations (73) *and* (74) *hold. In that event, any solution satisfying the robust CVaR constraint in the optimization problem* (13) *satisfies the corresponding true-distribution CVaR bound, and its true expected profit is no smaller than its robust objective guarantee.*

Remark 5.10 (Scope of the protection statement)**.** Theorem 5.9 is conditional on statistical coverage of both ambiguity sets. The controlled experiments below check the optimization structure and algorithmic behavior, but they neither estimate η and ζ from hotel data nor empirically assess the coverage probabilities in Equations (73)–(74). A deployment study would therefore require data-driven calibration of δ_j, Γ, and ε.

6. Computational experiments and real-world-calibrated illustration

6.1. Role of the experiments

The computational study has three linked roles. First, controlled cost-structured instances check the structural results, audit the certified cutting-plane implementation, and illustrate the cost policy implications of the model in a transparent reservation environment. Second, the scalability experiments evaluate the algorithmic efficiency of the finite robust counterpart through cutting-plane iterations, generated cuts, central processing unit (CPU) time, residuals, and finite-model optimality gaps. A matched formulation-level benchmark further compares delayed generation with the monolithic finite robust dual counterpart under the same LP backend, so the quality and CPU time are assessed without confounding the comparison by solver technology. Third, a public hotel booking data set is used only as a real-world-calibrated illustration of segment heterogeneity and calibration ranges.

The public booking records are not policy execution logs and therefore cannot be used to measure the computational efficiency of the optimization algorithm or to validate realized overbooking performance. We use them only to ground a real-world-calibrated illustration of segment heterogeneity and calibration ranges. Algorithmic efficiency is instead evaluated on controlled cost-structured instances, where the capacity, uncertainty sets, objective coefficients, active cuts, optimality gaps, and CVaR residuals are all directly observable.

The reported controlled optimization experiments use the continuous relaxation, which isolates the robust allocation structure and permits direct finite-model certification. Integer booking limits can be imposed in the same finite counterpart by declaring the booking variables integral, but integer versus continuous performance is not treated as a separate empirical contribution in this study.

The numerical analysis supports the analytical contribution by focusing on structural checks and optimization policy behavior rather than on a forecasting leaderboard. The main objects of interest are the direction of worst-case calibration error, phase-transition boundaries, the box-local diagnostic, the weighted budget active scenario certificate, the price of calibration uncertainty, uncertainty set sensitivity, finite convergence of Algorithm 1, and the calibration range grounding provided by the real-world-calibrated illustration. Accordingly, Section 7 reports the evidence in the same order as the claims it is intended to check: Reference law and critical fractile diagnostics, structural direction reversal checks, joint uncertainty policy response, finite-model certification and scalability, the matched formulation benchmark, and finally public data calibration grounding. This ordering is intended to make the theory–experiment link explicit rather than to present the computations as an undifferentiated simulation study.

6.2. Synthetic hotel generator and evaluation design

The controlled instances represent a hotel with capacity C and J reservation segments. Each segment is described by the primitive tuple

$$(N_j, a_j, r_j, f_j, \widehat{q}_j, \delta_j, \omega_j), \tag{75}$$

where the entries denote segment availability, room-equivalent load, realized stay value, retained cancellation value, nominal show-up probability, calibration half-width, and calibration budget weight. These primitives are chosen to create economically distinct segment types rather than to mimic a particular property's realized booking history. The benchmark show-up vector is then selected from the calibration set or from controlled misspecification regimes so that the direction reversal, budget allocation, and risk certificate claims can be audited under known data-generating settings.

The numerical values in these controlled instances are design parameters rather than estimates for a particular property. Their qualitative structure—heterogeneous cancellation exposure, segment-dependent booking value and load, and asymmetric consequences of overflow and unused capacity—is motivated by hotel booking control and overbooking research [12, 13, 33]. Quantities observable in the public booking records, such as segment shares, show-up rates, average daily rate (ADR) proxies, stay load proxies, and split-sample calibration bands, are used only for the real-world-calibrated illustration. Parameters not identified by those records, including overflow recovery cost, CVaR tolerance, and ambiguity radii, are therefore treated as controlled sensitivity parameters rather than externally estimated constants.

The controlled study uses a standardized lower-truncated skew-normal law as the baseline residual reference because Section 7.1 explicitly audits its fit and the associated asymmetric loss geometry. The main uncertainty comparison then varies the Wasserstein radius around this common reference law. Alternative residual families are not used as separate main-text performance benchmarks, so the reported policy comparisons remain tied directly to the distributional model and finite reformulation analyzed in Section 5.

The numerical analysis is organized around the structural claims rather than around a separate checklist of named experiments. Complete enumeration in small instances checks the endpoint rule, the extreme-point adversary, and the weighted-budget ranking. Phase-transition maps vary overflow cost and demand pressure to check the transition condition $g = \tau_j$. Local diagnostics report $M_j^{\text{box}}(g)$ only as coordinatewise summaries, while the global weighted budget certificate is evaluated from the active-scenario condition in Theorem 4.12. The joint uncertainty analysis varies $\Gamma \in \{0, 1, 2, 3, 4\}$ and $\varepsilon \in \{0, 0.025, 0.05, 0.075, 0.10, 0.15, 0.20\}$, yielding 35 controlled cases that report robust value, PCU, expected service load, portfolio composition, overflow CVaR, feasibility, and the finite-model optimality gap. The scalability study varies $J \in \{8, 16, 32, 64, 96\}$ and reports runtime, generated cuts, cutting-plane iterations, residuals, and optimality gaps. A separate matched benchmark at $J \in \{8, 16, 32\}$ compares the delayed generation formulation with the complete monolithic finite counterpart using the same LP backend and common numerical resolution. Finally, the public booking records are used only to calibrate segment shares, cancellation rates, show-up rates, ADR proxies, stay load proxies, and split-sample calibration bands. This public data illustration grounds parameter ranges for the optimization study, but it is not an evaluation of an implemented booking policy.

Table 3 summarizes the links between the structural results and the numerical diagnostics, so a separate synthetic design table is not repeated here.

6.3. Optimization policy benchmarks

The computational comparison is restricted to optimization policy benchmarks. The nominal policy uses the empirical show-up distribution $q = \widehat{q}$ without Wasserstein ambiguity, that is, $\varepsilon = 0$. The calibration-robust policy keeps $\varepsilon = 0$ but allows the show-up distribution to vary within the calibration uncertainty set $\mathcal{U}_{\mathrm{cal}}(\Gamma)$. The distributionally robust policy fixes $q = \widehat{q}$ while introducing a positive Wasserstein radius $\varepsilon > 0$, thereby isolating the effect of aggregate load distributional ambiguity. The joint robust policy combines both forms of uncertainty and evaluates the resulting booking decision under the robust CVaR certificate. These four uncertainty specifications are the policy comparisons used in the main results; no nonimplementable oracle policy is used to support the reported conclusions.

This benchmark design isolates the optimization value of calibration protection, distributional robustness, and their joint use. It does not include a machine learning cancellation prediction model, area under the receiver operating characteristic curve (AUC) comparison, calibration curve, Shapley additive explanations (SHAP) analysis, or empirical hotel-profit estimate. These exclusions are deliberate: the purpose of the experiments is to evaluate the proposed robust booking control model and its uncertainty set specifications, not to compare predictive classifiers or to claim realized hotel-profit effects.

7. Numerical results

The main computational study uses controlled synthetic data. Its role is to audit the structural claims and the certified finite optimization model, not to establish realized performance superiority for a particular hotel. All monetary quantities in the controlled optimization experiments are synthetic units. Within each comparison, policies use the same controlled instance primitives and the same deterministic finite reference approximation; thus the differences are attributable to the uncertainty specification or formulation being compared rather than to different simulated samples. A separate real-world-calibrated illustration based on public booking records is reported at the end of this section to ground the segment-level cancellation and load parameters; it is used only for calibration range grounding.

7.1. Aggregate-load reference law and asymmetric capacity loss structure

Before analyzing portfolio responses, we first calibrate the one-dimensional aggregate load law that underlies the robust loss term and the numerical critical fractile calculations. For each arrival day t, let $\widehat{L}_t$ denote the nominal expected show-up load and let L_t denote the realized show-up load. The disturbance proxy is constructed as

$$z_t = \frac{(L_t - \widehat{L}_t)/\sqrt{\max\{\widehat{L}_t, 1\}} - \mu_z}{\sigma_z}, \tag{76}$$

where μ_z and σ_z are the sample mean and sample standard deviation of the variance-stabilized residual series. The resulting standardized series contains 1,096 synthetic daily observations spanning 2023–2025. Its empirical skewness is -0.480, and its empirical excess kurtosis is 0.315.

Figure 2 summarizes the fitted reference laws and the induced asymmetric capacity loss geometry. Figure 2(a) shows that the synthetic daily residuals are negatively skewed, which motivates an asymmetric reference law. Figure 2(b) compares the normal, skew-normal, and lower-truncated skew-normal fits. The normal benchmark is visibly too symmetric, while the two asymmetric laws track the empirical shape closely. Figure 2(c) confirms that both asymmetric fits improve substantially on the normal law in the operationally important tails. Figure 2(d) reports the expected normalized loss under alternative overflow-to-vacancy cost ratios. For the illustrative normalized-loss ratio $c_o/c_u = 4$, the theoretical critical fractile target implied by Corollary 4.2 is $d^\star = -0.851$, while the simulated loss-minimizing displacement is $d^\star = -0.867$. The numerical gap is only 0.016 standardized load units, providing a close numerical check of the analytical target rule in Equation (21). This unit-free diagnostic is separate from the booking control baseline $c_o = 420$ and $c_u = 65$ used in the portfolio experiments below.

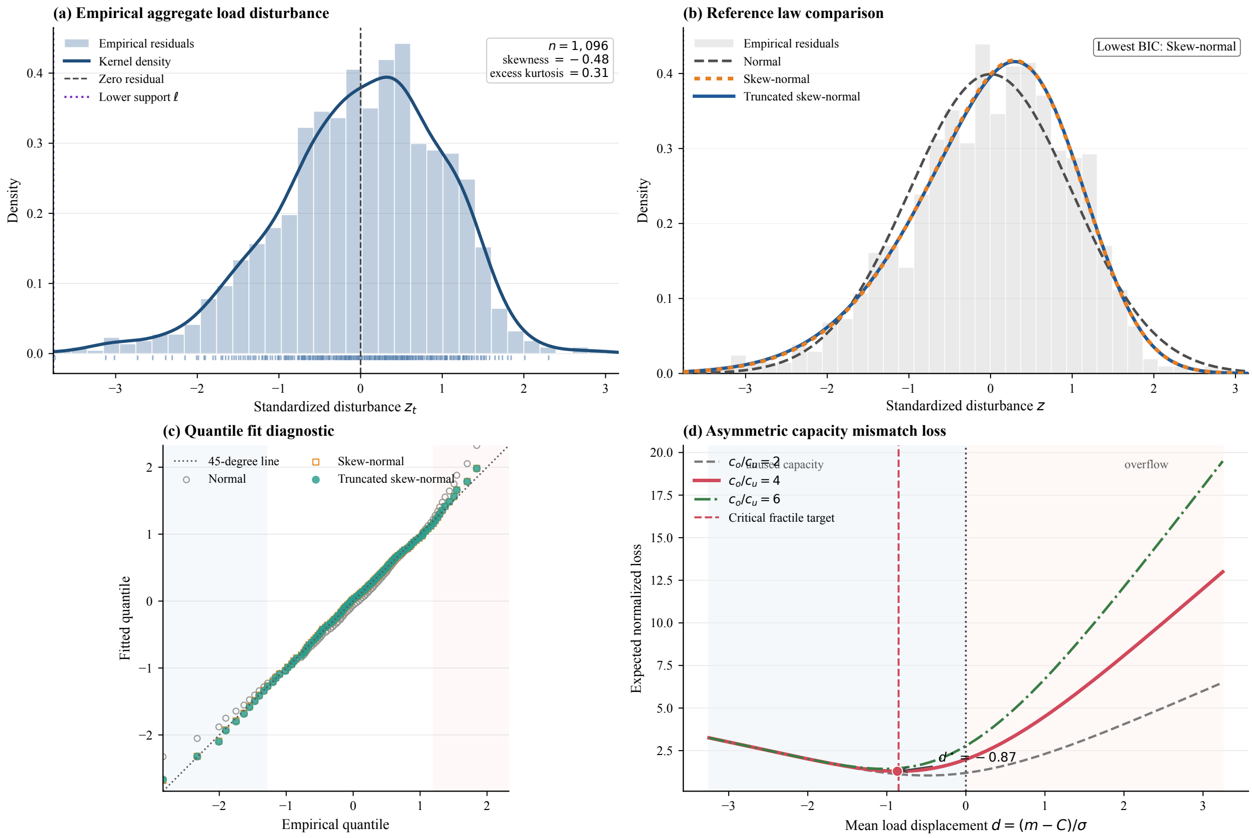

Figure 2. Aggregate load uncertainty and asymmetric capacity loss structure. (a) Empirical distribution of the standardized aggregate-load disturbance together with the kernel estimate and estimated lower support. (b) Fitted normal, skew-normal, and lower-truncated skew-normal reference laws. (c) Empirical fitted quantile diagnostic, with emphasis on the lower and upper tails. (d) Expected normalized capacity mismatch loss under alternative overflow-to-vacancy cost ratios; the highlighted point identifies the base case loss-minimizing displacement.

Table 4 reports the reference law diagnostics for the standardized aggregate load residuals. Relative to the normal law, both asymmetric specifications reduce the Akaike information criterion (AIC) by more than 41 points and approximately halve the empirical Kolmogorov–Smirnov (KS) discrepancy.

The skew-normal specification achieves the lowest Bayesian information criterion (BIC), whereas the truncated skew-normal specification attains the lowest AIC and the smallest KS discrepancy. We therefore retain the truncated skew-normal law as the baseline reference distribution because its empirical fit is competitive and its finite lower support is structurally consistent with a bounded lower tail in realized service load. Its fitted shape, location, scale, and lower support parameters are -2.196333, 1.059420, 1.465088, and $\ell = -3.763523$, respectively.

The KS discrepancy is used only as an empirical distance measure and is not interpreted as a classical goodness-of-fit test, because the model parameters are estimated from the same sample. In Figure 2(d), the illustrative normalized ratio is $c_o/c_u = 4$, which implies a critical fractile of 0.8, an analytical target $d^\star = -0.851$, and a simulated minimum $d^\star = -0.867$. It should not be confused with the monetary cost pair $(c_o, c_u) = (420, 65)$ used in the subsequent booking control experiments.

Table 4. Reference law fit for standardized aggregate load residuals.

Distribution	Log-likelihood	AIC	BIC	KS discrepancy	ΔAIC	ΔBIC
Normal	-1554.656	3113.313	3123.312	0.04038	41.356	36.302
Skew-normal	-1533.006	3072.011	**3087.010**	0.01802	0.054	**0.000**
Truncated skew-normal	-1531.978	**3071.957**	3091.955	**0.01754**	**0.000**	4.945

7.2. Synthetic segment design

Table 5 reports the baseline reservation segments. The design deliberately combines high- and low-value segments, narrow and wide calibration intervals, and single- and multi-room-equivalent loads. This creates both revenue-dominant and congestion-dominant regimes.

Table 5. Baseline synthetic hotel segments.

Segment	N_j	a_j	r_j	f_j	$\widehat{q}_j$	δ_j	ω_j	τ_j
Advance saver	24	1.0	145	35	0.82	0.035	1.0	110.0
Flexible retail	22	1.0	175	10	0.72	0.075	1.2	165.0
Corporate	18	1.0	190	20	0.91	0.030	0.9	170.0
OTA leisure	26	1.0	155	5	0.68	0.085	1.3	150.0
Group	14	1.6	260	55	0.86	0.055	1.1	128.1
Premium	12	1.0	240	25	0.94	0.025	0.8	215.0
Weekend package	20	1.3	210	45	0.79	0.050	1.0	126.9
Loyalty	16	1.0	185	15	0.89	0.040	0.9	170.0

7.3. Structural checks

Figure 3 provides a direct numerical audit of Theorems 4.7 and 4.8. The marginal capacity price is evaluated from Equation (18), while each segment threshold is recomputed from Equation (28). The maximum discrepancy between the recomputed and stored threshold values is zero to numerical precision. Under the baseline environment $C = 100$, $\sigma = 8$, $c_u = 65$, $c_o = 420$, and $m = 98$, Equation (18) gives

$$g(98) = 145.857277. \tag{77}$$

Figure 3(a) plots the segment-specific economic sensitivity

$$\psi_j(g) = (r_j - f_j) - a_j g = a_j(\tau_j - g), \tag{78}$$

which is the coefficient governing the derivative in Equation (29). The zero of each line occurs exactly at $g = \tau_j$. Consistent with Equation (31), $\psi_j(g) > 0$ implies a locally adverse downward perturbation toward $\widehat{q}_j - \delta_j$, whereas $\psi_j(g) < 0$ implies a locally adverse upward perturbation toward $\widehat{q}_j + \delta_j$. At the baseline price in Equation (77), advance saver, weekend package, and group select the upper endpoint, while online travel agency (OTA) leisure, flexible retail, corporate, loyalty, and premium select the lower endpoint.

Figure 3(b) evaluates the phase-transition condition in Equation (32) for the flexible retail segment, whose threshold is $\tau_j = 165$. The black contour is the exact solution set $g(m, c_o, c_u) = 165$. At the baseline overflow cost $c_o = 420$, the contour crosses the mean load axis at

$$m^{\text{phase}} = 98.763518. \tag{79}$$

Because the baseline mean load is $m = 98 < m^{\text{phase}}$, the baseline operating point lies in the lower endpoint region, in agreement with the sign of $\psi_j(145.857277) = 19.142723$. Increasing c_o shifts the phase boundary toward lower mean loads: The boundary decreases from 113.386 at $c_o = 180$ to 94.073 at $c_o = 900$. Thus, more severe overflow consequences cause upward show-up probability error to become adverse at progressively lower utilization levels.

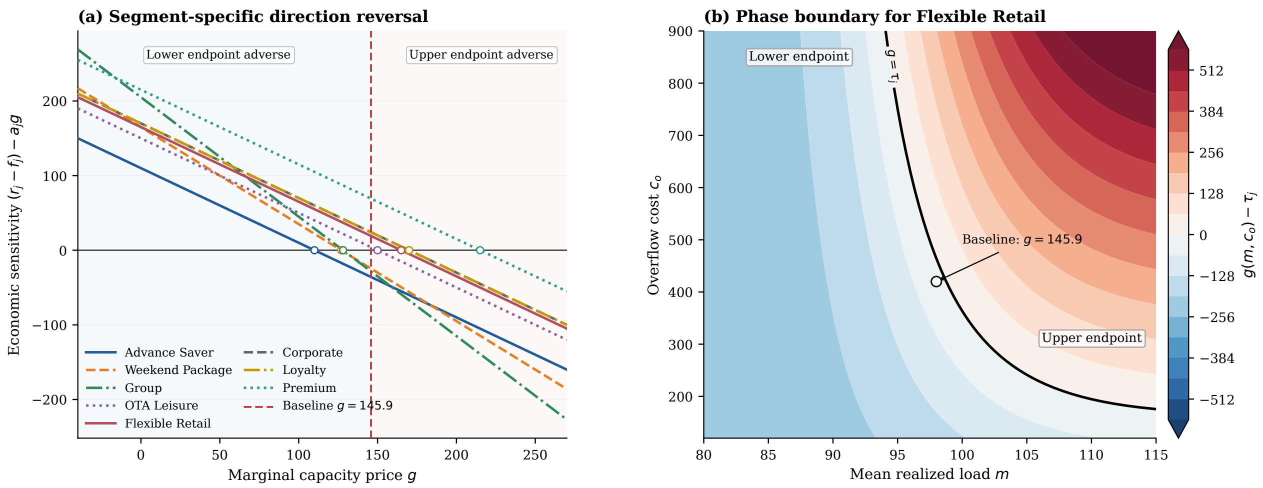

Figure 3. Calibration direction reversal and phase transition. (a) Segment specific sensitivity $\psi_j(g)$ in Equation (78); each zero marker occurs at the threshold τ_j defined in Equation (28), and the dashed vertical line marks the baseline value $g = 145.857277$. (b) Values of $g(m, c_o, c_u) - \tau_j$ for flexible retail. The black zero contour is the phase boundary in Equation (32), and the open marker identifies the baseline operating point $(m, c_o) = (98, 420)$.

Table 6 reports the segment-level baseline diagnostics. The sign of $g - \tau_j$ and the sign of $\psi_j(g)$ are opposite by Equation (78); the resulting local adverse directions coincide with Theorem 4.7.

Table 6. Baseline direction-reversal diagnostics.

Segment	τ_j	$g-\tau_j$	$\psi_j(g)$	q_j^{wc}	Adverse calibration endpoint
Advance saver	110.000	35.857	-35.857	0.855	Upper: $\widehat{q}_j+\delta_j$
Weekend package	126.923	18.934	-24.614	0.840	Upper: $\widehat{q}_j+\delta_j$
Group	128.125	17.732	-28.372	0.915	Upper: $\widehat{q}_j+\delta_j$
OTA leisure	150.000	-4.143	4.143	0.595	Lower: $\widehat{q}_j-\delta_j$
Flexible retail	165.000	-19.143	19.143	0.645	Lower: $\widehat{q}_j-\delta_j$
Corporate	170.000	-24.143	24.143	0.880	Lower: $\widehat{q}_j-\delta_j$
Loyalty	170.000	-24.143	24.143	0.850	Lower: $\widehat{q}_j-\delta_j$
Premium	215.000	-69.143	69.143	0.915	Lower: $\widehat{q}_j-\delta_j$

Notes: The baseline marginal capacity price is $g = 145.857277$. The sensitivity is $\psi_j(g) = (r_j - f_j) - a_j g$. A negative value gives a locally adverse upward direction and a positive value gives a locally adverse downward direction, as stated in equation (31). Values are rounded for presentation.

Figure 4 reports two complementary diagnostics at the canonical weighted budget solution for $\Gamma = 3$ and $\varepsilon = 0$. Figure 4(a) shows the box-local coordinatewise quantity in Equation (37). It is presented only as an interpretive benchmark and not as a global KKT certificate for the weighted budget model. Figure 4(b) reports the local continuous knapsack ranking in Proposition 4.10. The exact global certificate is instead evaluated from the active worst-case extreme points using Theorem 4.12.

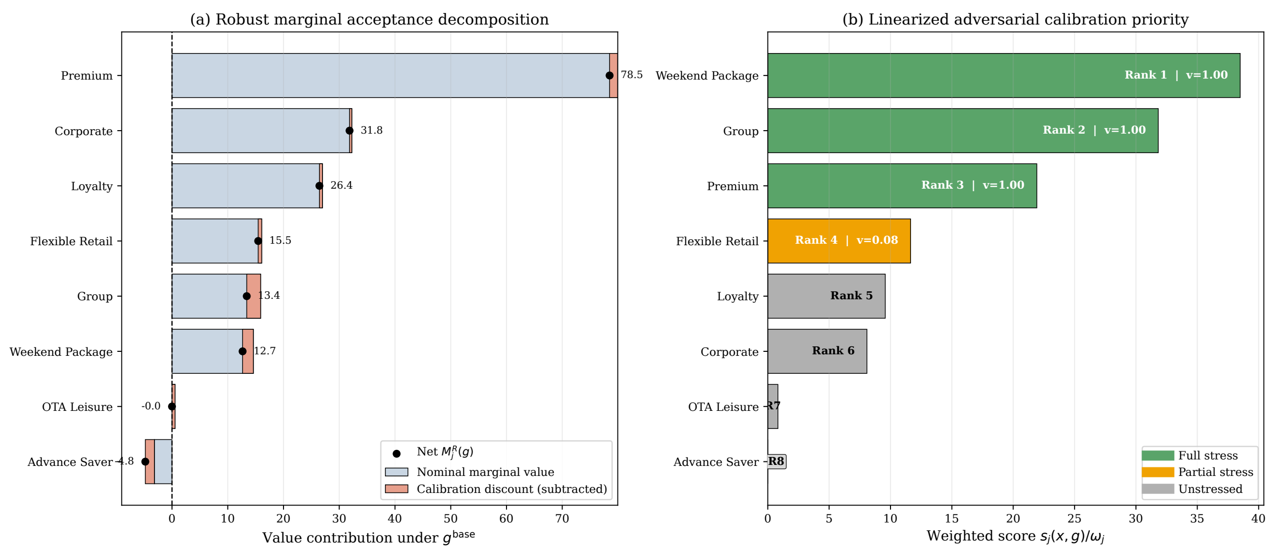

Figure 4. Local calibration diagnostics at the canonical weighted budget solution. (a) Decomposition of the box-local diagnostic $M_j^{\mathrm{box}}(g)$ into nominal marginal value and the coordinatewise calibration discount. (b) Linearized adversarial priority $s_j(x,g)/\omega_j$. These diagnostics are interpretive local summaries rather than a global weighted budget optimality certificate.

7.4. *Joint uncertainty price and portfolio response*

The joint uncertainty analysis evaluates the complete grid

$$\Gamma \in \{0, 1, 2, 3, 4\}, \qquad \varepsilon \in \{0, 0.025, 0.05, 0.075, 0.10, 0.15, 0.20\}, \tag{80}$$

using the same finite piecewise-linear reference approximation and certified cutting-plane solver for every case. This common computational basis is important: It prevents the uncertainty price comparisons from mixing objective definitions, local optimizers, or inconsistent risk approximations. All 35 grid points are feasible and terminate with an optimal master solution. The maximum final value gap is 3.64×10^{-12}, the maximum CVaR violation is 2.84×10^{-14}, and the maximum support violation is zero.

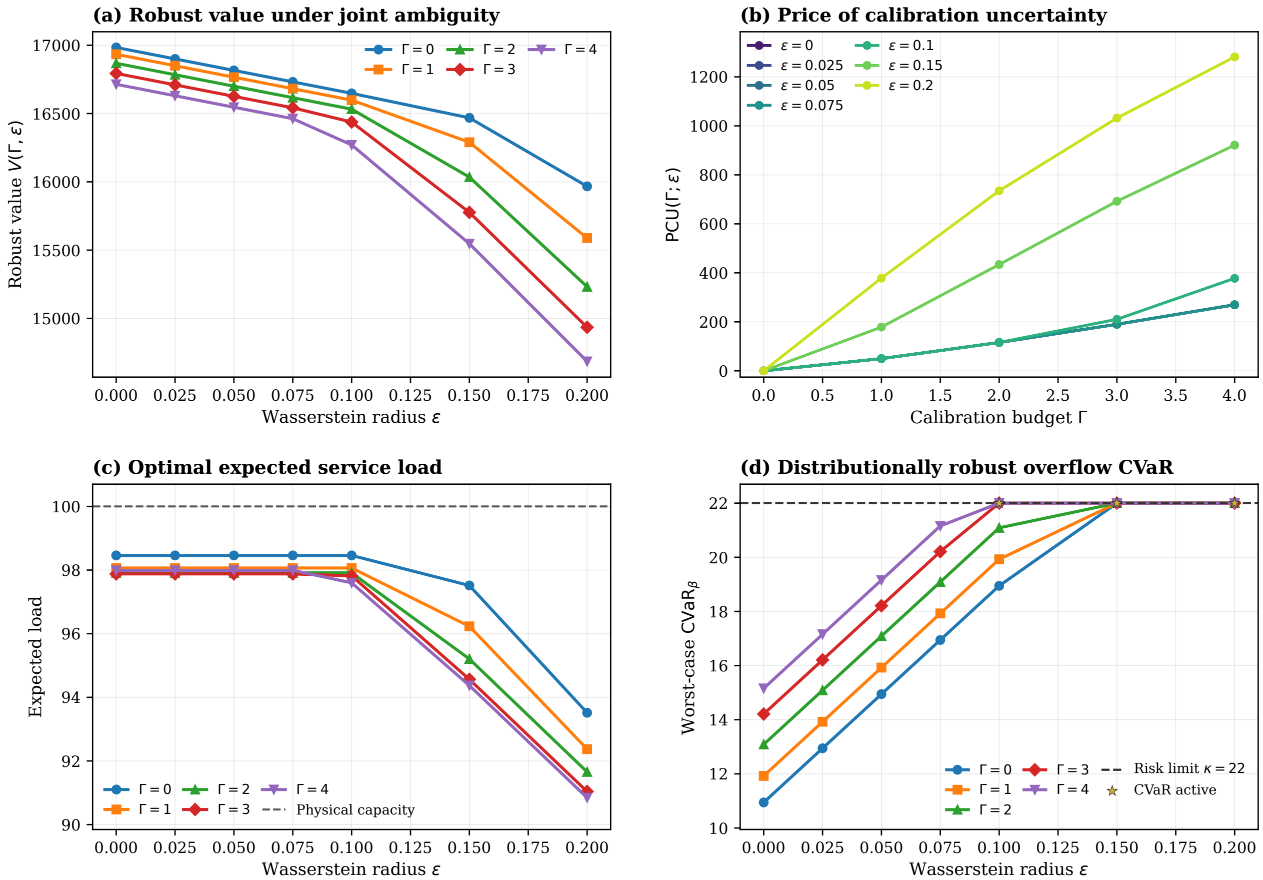

Figure 5. Joint effects of calibration and distributional ambiguity. (a) Robust value over the Wasserstein radius for each calibration budget. (b) Price of calibration uncertainty reported directly in objective value units, without normalization by a fixed portfolio sensitivity bound. (c) Nominal expected service load, with physical capacity shown for reference. (d) Worst-case overflow CVaR, identifying cases in which the risk limit $\kappa = 22$ is active.

Figure 5 reports the four principal comparative statics. Figure 5(a) confirms that robust value is nonincreasing in both ambiguity parameters. Figure 5(b) shows that the price of calibration uncertainty is non-negative and nondecreasing in the calibration budget, consistent with the global monotonicity statement in Theorem 4.13. The analytical upper bound in that theorem is a fixed-portfolio sensitivity bound and is therefore not used here as a global optimal value certificate. Figure 5(c) identifies the policy adjustment channel: Expected service load is nearly unchanged for small radii but decreases once the robust CVaR constraint becomes active. Figure 5(d) shows that larger calibration budgets cause the CVaR limit to bind at smaller Wasserstein radii. Thus, calibration error and residual law ambiguity are not interchangeable. Calibration ambiguity changes which segments are attractive, whereas Wasserstein ambiguity raises aggregate tail protection and eventually activates the service risk constraint.

Table 7 isolates four representative policies. Under calibration ambiguity alone, the robust value falls by 189.997 units and the expected service load falls from 98.460 to 97.879. Under distributional ambiguity alone, the portfolio is unchanged at the canonical radius $\varepsilon = 0.1$, while the robust value falls by exactly 336 units and overflow CVaR increases from 10.945 to 18.945. This exact difference follows from the fixed-scale Wasserstein surcharge in Proposition 5.3. Under joint ambiguity, the robust value is 16,437.585, the total protection cost is 546.269, and the CVaR constraint is bounded at 22. The joint policy therefore cannot be reconstructed by simply adding two scalar penalties to the nominal policy: It also changes the accepted portfolio and creates two boundary segments.

The distributional-only differences also provide a direct numerical check of Proposition 5.3: $\varepsilon\sigma c_o = 0.1 \times 8 \times 420 = 336$ and $\varepsilon\sigma/(1-\beta) = 0.1 \times 8/0.1 = 8$. In contrast, calibration ambiguity changes accepted quantities. The four-policy comparison therefore gives a transparent computational example of why Γ and ε are distinct uncertainty controls rather than interchangeable conservatism parameters.

The same comparison also shows a modest interaction under joint protection. Calibration-only protection costs 189.997 units at $\varepsilon = 0$, whereas the calibration component conditional on $\varepsilon = 0.1$ is $16{,}647.854 - 16{,}437.585 = 210.269$ units. Hence the joint protection cost 546.269 exceeds the mechanical sum $189.997 + 336$ by 20.272 units because the portfolio is reoptimized under the joint uncertainty set and the robust CVaR constraint becomes binding. This is an instance-specific interaction, not a general claim of superadditivity.

Table 7. Representative policies under calibration and distributional ambiguity.

Policy	Γ	ε	Robust value	Protection cost	Expected load	Overflow CVaR	Boundary segment(s)
Nominal	0	0	16,983.854	0.000	98.460	10.945	OTA leisure
Calibration robust	3	0	16,793.857	189.997	97.879	14.210	OTA leisure
Distributionally robust	0	0.1	16,647.854	336.000	98.460	18.945	OTA leisure
Joint robust	3	0.1	16,437.585	546.269	97.818	22.000	Flexible retail; OTA leisure

Notes: Protection cost is measured relative to the nominal robust value. Monetary values are synthetic units. The joint-policy calibration component is approximately 210.269 units after accounting for the 336-unit distributional surcharge. All four policies are feasible in the certified finite model.

Figure 6 makes the portfolio channel explicit. At $\varepsilon = 0$, increasing Γ primarily lowers the worst-case value and causes only limited portfolio movement. At $\varepsilon = 0.1$, the risk constraint begins to influence the composition for larger budgets. At $\varepsilon = 0.2$, the policy response is visibly segment specific: Total accepted bookings fall and the allocation shifts across reservation classes rather than applying a uniform booking discount. This is the operational counterpart of the weighted budget adversary in Proposition 4.10.

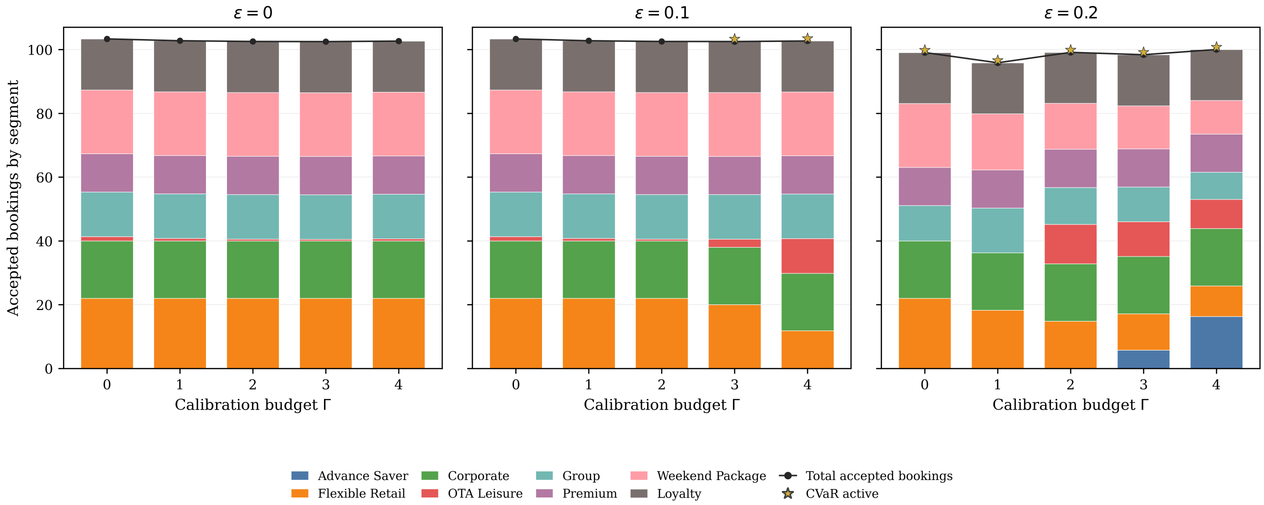

Figure 6. Optimal synthetic booking portfolio composition under joint ambiguity. The three displayed cases report segment-level accepted quantities across calibration budgets at $\varepsilon \in \{0, 0.1, 0.2\}$. Stars identify solutions for which the robust overflow CVaR constraint is active. The segment-specific adjustments show that calibration ambiguity changes the portfolio's composition rather than imposing a uniform reduction in accepted bookings.

The composition patterns in Figure 6 also distinguish three operational regimes. With little distributional ambiguity, calibration protection mainly reallocates bookings among segments while leaving aggregate acceptance relatively stable. As the Wasserstein radius increases, the robust tail constraint becomes more influential and the policy begins to reduce total accepted load as well as redistribute it. Once the CVaR boundary is reached, additional calibration protection is absorbed selectively by the segments whose value-to-load and uncertainty profiles make them most costly under stress. This progression is consistent with the analytical separation between segment-level calibration ambiguity and aggregate residual-load ambiguity: the former primarily changes portfolio composition, whereas the latter increasingly governs the amount of capacity protection required.

7.5. Uncertainty-grid policy implications

The complete 35-case grid provides policy information beyond the four representative cases in Table 7. Table 8 summarizes the distinct operational roles of calibration ambiguity, distributional ambiguity, and their interaction. Calibration ambiguity primarily changes segments' attractiveness and the portfolio's composition, whereas Wasserstein ambiguity raises aggregate tail protection and can activate the robust CVaR constraint. Their joint effect is therefore not equivalent to adding two independent scalar penalties. Across all grid points, robust value is nonincreasing in both ambiguity parameters and the price of calibration uncertainty is non-negative, consistent with Theorems 4.15 and 4.13. All finite-model solutions satisfy the robust CVaR, support, and optimality checks within the numerical tolerance of 10^{-8}. Detailed numerical certification diagnostics are reported in Appendix B, Table A1.

Table 8. Policy interpretation of the 35-case uncertainty grid.

Uncertainty dimension	Grid evidence	Booking control implication
Calibration budget Γ	Robust value is nonincreasing in Γ and the price of calibration uncertainty is non-negative and nondecreasing over the tested budgets.	Protection is directed toward segment-level probability misspecification and may change the relative attractiveness and accepted quantities of individual segments.
Wasserstein radius ε	Robust value is nonincreasing in ε; worst-case overflow CVaR rises with distributional protection and eventually reaches the risk limit.	Protection is applied to aggregate residual load uncertainty and increasingly emphasizes tail-risk control rather than reservation-level probability adjustment.
Joint ambiguity (Γ, ε)	Larger calibration budgets cause the robust CVaR constraint to bind at smaller Wasserstein radii, and expected service load declines once the risk constraint becomes active.	Calibration and distributional ambiguity interact through both portfolio composition and capacity protection; the joint policy cannot be reconstructed from two independent scalar surcharges.
Finite-model audit	All 35 cases are feasible and solved optimally; the maximum CVaR residual is 2.84×10^{-14} and the maximum final value gap is 3.64×10^{-12}.	The reported comparative statics are numerically consistent with the finite quadrature and piecewise-linear model used in computation.

Note: The audit applies to the finite quadrature and piecewise-linear model solved computationally and is not claimed as an exact certificate for an un-discretized infinite-dimensional problem.

Table 9 gives a compact numerical response to the two computational questions raised by the revision: Whether the proposed algorithm is efficient on the tested instances and whether the uncertainty set specification affects the results. Panel A summarizes the measured scaling runs without repeating nearly identical rows. Panel B reports selected nonredundant uncertainty-set specifications, separating calibration budget effects from Wasserstein radius effects. The direct solution quality and CPU comparison with the monolithic robust dual formulation is reported separately in Table 10, because it uses a common benchmark resolution chosen specifically for the matched formulation comparison.

Table 9. Computational efficiency and uncertainty-set sensitivity.

Panel A. Algorithmic efficiency

Instance range	Runs solved	Iterations	Generated cuts	Time (s)	Certification
J = 8–96	25/25	6–8	9–11	0.087–0.099	zero gap/residuals

Panel B. Uncertainty set sensitivity

Spec.	Γ	ε	Robust value	Accepted bookings	Overflow CVaR	PCU
$\varepsilon = 0$	0.00	0.000	16,983.85	103.35	10.95	0.00
$\varepsilon = 0$	1.00	0.000	16,934.04	102.77	11.92	49.82
$\varepsilon = 0$	2.00	0.000	16,868.14	102.54	13.09	115.72
$\varepsilon = 0$	3.00	0.000	16,793.86	102.49	14.21	190.00
$\varepsilon = 0$	4.00	0.000	16,714.17	102.66	15.16	269.68
$\Gamma = 3$	3.00	0.050	16,625.86	102.49	18.21	190.00
$\Gamma = 3$	3.00	0.100	16,437.59	102.52	22.00	210.27
$\Gamma = 3$	3.00	0.200	14,934.95	98.38	22.00	1,031.84

Notes: Panel A summarizes all successful scalability runs over the tested instance size range. Panel B reports selected nonredundant uncertainty-set specifications. PCU denotes the price of calibration uncertainty. The public data analysis is reported separately as a real-world-calibrated illustration, not as evidence of implemented policy performance.

7.6. *Finite convergence and scalability*

The scalability analysis evaluates the computational behavior of the certified cutting-plane implementation as the segment dimension increases from $J = 8$ to $J = 96$. For each size, one warm-up run is followed by five measured runs in independent worker processes. End-to-end time includes model construction, reference piecewise-linear approximation, and the cutting-plane optimization loop. Peak memory is measured by the unique set size (USS), avoiding the cumulative high-water-mark artifact that arises when all cases are measured inside a single long-lived notebook process.

Figure 7 and Panel A of Table 9 show stable computational effort over the tested range. The median iteration count remains between 6 and 8, and the median number of generated cuts remains between 9 and 11. Median end-to-end time remains below 0.10 seconds, while median peak USS remains close to 40 megabytes (MB). All 25 measured runs terminate optimally; the largest observed gap is 3.64×10^{-12}, and no CVaR or support violation is observed. These results do not establish constant-time complexity. They show that, for the tested synthetic range and fixed numerical resolution, delayed generation avoids growth in the number of active scenarios and pieces.

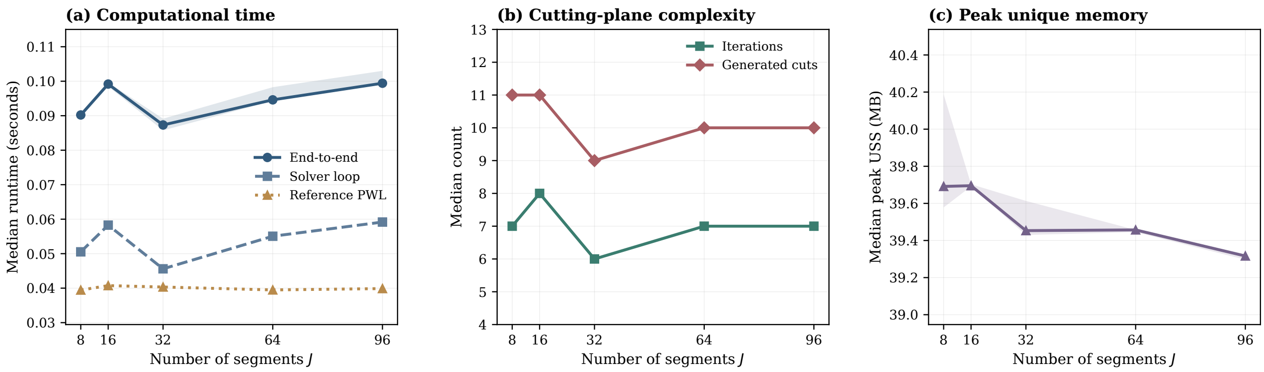

Figure 7. Computational scalability of the certified cutting-plane algorithm. (a) Median reference piecewise-linear (PWL), solver-loop, and end-to-end time over five independent measured runs; the shaded band is the interquartile range for end-to-end time. (b) Median cutting-plane iterations and generated cuts. (c) Median peak USS memory with its interquartile range.

7.7. *Matched benchmark against the monolithic finite robust-dual formulation*

To address solution quality and CPU time directly, we compare the delayed-generation implementation with the complete monolithic finite robust-dual counterpart while holding the optimization backend fixed. Both formulations use the same SciPy optimization environment with the HiGHS linear programming solver (SciPy/HiGHS), the same continuous decision model, the same baseline uncertainty setting $(\Gamma, \varepsilon) = (3, 0.10)$, and the same finite reference approximation. The comparison is therefore formulation-level rather than solver-brand-level: it isolates the effect of delayed constraint generation from differences in optimization software.

For this matched benchmark, both formulations use 3001 deterministic reference quantiles and a 1001-point load grid. This common lower numerical resolution keeps the monolithic counterpart tractable and is used only for the matched formulation comparison; it does not replace the higher-resolution policy grid reported in the preceding subsections. For each $J \in \{8, 16, 32\}$, one warm-up run is followed by five measured runs, and the execution order is alternated across repetitions. Table 10 reports method-only solution time after the common reference construction, because the reference approximation is identical for both formulations.

The benchmark is intentionally conservative in attributing computational gains. Reference-law construction, data preparation, and common numerical preprocessing are excluded from the timing ratio, so the comparison measures only the additional optimization work required by each formulation after identical inputs have been created. The delayed-generation method is not given a different stopping rule or feasibility tolerance, and both formulations are evaluated against the same robust-value, CVaR, and support checks. Consequently, a smaller solution time is informative only when the resulting portfolio and objective value remain numerically indistinguishable from those of the complete finite counterpart.

Table 10. Matched delayed generation and monolithic finite robust dual benchmark.

J	Delayed CPU (s)	Monolithic CPU (s)	Time ratio	Generated cuts	Full constraints	Full variables	Max. $\|\Delta V\|$
8	0.0456	0.0757	1.66	11	18,009	18,018	7.13×10^{-10}
16	0.0539	0.1211	2.25	11	34,017	34,034	6.91×10^{-11}
32	0.0415	0.2401	5.78	9	66,033	66,066	0

Notes: CPU entries are the medians over five measured runs after one warm-up run. Both formulations use the same SciPy/HiGHS backend and the same finite reference approximation. The maximum segment-level portfolio difference over all matched runs is 9.20×10^{-10}. Every matched run satisfies the robust value, CVaR, and support certification tolerances.

Two features of Table 10 are especially relevant for interpreting computational efficiency. First, the portfolio and objective discrepancies remain at numerical tolerance for every tested size, so the timing comparison does not trade solution quality for speed. Second, the number of active cuts remains essentially flat even as the complete finite model expands from about 18,000 to more than 66,000 constraints. The increasing time ratio is therefore associated with avoiding inactive robust constraints rather than with a different optimizer or a coarser decision model. Because both formulations share the same finite reference construction, these results isolate the formulation-level benefit of generating only constraints that become relevant at the incumbent solution.

The two formulations reproduce the same certified optimum to numerical precision while instantiating very different model sizes. Delayed generation uses only 9–11 active cuts, whereas the monolithic formulation contains 18,009–66,033 constraints over the tested sizes. The relative CPU advantage increases from 1.66 at $J = 8$ to 5.78 at $J = 32$. These results do not establish an asymptotic complexity rate; they show that, in the matched instances, delayed generation preserves the solution's quality while avoiding the a priori construction of most finite robust constraints.

Figure 8 consolidates the formulation-level benchmark and the canonical uncertainty response in a single audit view. Figure 8(a), (b) uses the matched delayed generation versus monolithic experiment reported in Table 10; Figure 8(c), (d) uses the canonical nominal, calibration-only, distributional-only, and joint robust cases from the reported uncertainty grid. Thus, the figure introduces no separate illustrative values: Every plotted quantity is computed from the numerical experiments reported in this section.

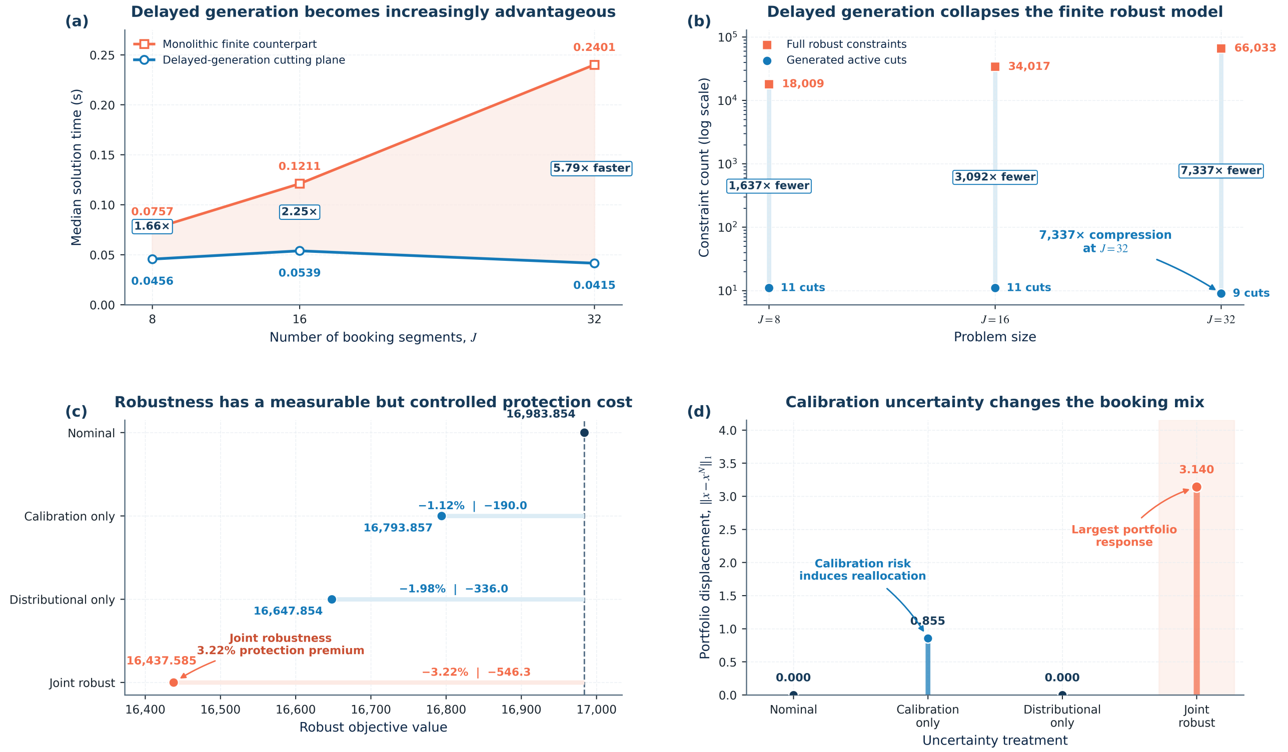

Figure 8. Integrated computational and uncertainty-response summary. (a) Median method-only solution time for the delayed-generation cutting-plane formulation and the matched monolithic finite robust dual counterpart under the same SciPy/HiGHS backend and finite reference approximation. (b) Full finite-model constraint count versus the active cuts generated by delayed generation, shown on a logarithmic scale. (c) Robust objective value under the nominal, calibration-only, distributional-only, and joint-robust specifications, with the objective reduction measured relative to the nominal policy. (d) ℓ_1 displacement of each optimized booking portfolio from the nominal portfolio. Figure 8(a), (b) reports the matched formulation benchmark at $(\Gamma, \varepsilon) = (3, 0.10)$, whereas Figure 8(c), (d) reports the canonical four-policy comparison from the uncertainty grid experiment.

7.8. *Real-world-calibrated illustration*

The public data illustration is included to ground the synthetic parameter ranges in observable booking patterns, not to evaluate the robust policy as an implemented operating decision. The illustration uses a public hotel booking data set hosted on Kaggle and described by its provider as an updated and cleaned version of the original Hotel Booking Demand data set [48]. The original data source is documented by Antonio et al. [49]. We use the public records only to summarize segment frequencies, cancellation frequencies, show-up rates, ADR proxies, and stay load proxies. The records do not contain the information required for a realized overbooking policy evaluation, such as true room capacity decisions, rejected demand, walking costs, or realized booking control actions. For this reason, the controlled cost-structured experiments remain the main computational evidence, while the public data illustration is used only to show that the model's heterogeneous cancellation and load primitives are operationally plausible.

Table 11 reports the five representative public data segments retained for the main text. In the segment labels, TA denotes travel agent and TO denotes tour operator. The segments display substantial heterogeneity in show-up behavior: Groups have a show-up rate of 0.389, whereas direct and corporate bookings have show-up rates of 0.847 and 0.813, respectively. The table also reports segment-level booking shares, ADR proxies, stay load proxies, nominal training period show-up rates, and calibration bands. These quantities map naturally to the optimization primitives $(N_j, \widehat{q}_j, r_j, a_j, \Delta_j)$, but they are not used to claim realized profit gains.

Table 11. Real-world-calibrated illustration of segment heterogeneity from public hotel booking records.

Segment	N	Share	Cancel	Show-up	ADR	Load	$\widehat{p}_j$	Δ_j	Lower	Upper
Online TA	56,477	0.473	0.367	0.633	109.8	3.58	0.631	0.020	0.611	0.651
Offline TA/TO	24,219	0.203	0.343	0.657	85.5	3.91	0.659	0.024	0.635	0.683
Groups	19,811	0.166	0.611	0.389	70.0	2.99	0.390	0.020	0.370	0.410
Direct	12,606	0.106	0.153	0.847	105.0	3.22	0.851	0.036	0.815	0.887
Corporate	5,295	0.044	0.187	0.813	65.0	2.10	0.814	0.038	0.776	0.852

Notes: The table is a real-world-calibrated illustration, not implemented policy evidence. TA denotes travel agent and TO denotes tour operator. Share is the empirical booking share. Cancel and Show-up are the empirical cancellation and show-up frequencies. ADR is the median ADR proxy, Load is the mean stay-load proxy, $\widehat{p}_j$ is the training-period nominal show-up rate, and Δ_j is the split-sample calibration band. The public records do not contain rejected demand, walking costs, realized overbooking-control actions, or true capacity decisions.

A split-sample audit, reported in Appendix C, shows that held-out show-up rates fall within the constructed calibration bands for all five representative segments. This supports the use of the public records for calibration-range illustration, but it should not be interpreted as evidence of implemented booking-policy performance. Because the held-out observations come from the same public data source and the audit evaluates only segment-level show-up-rate coverage, we describe it as a held-out calibration audit rather than as external validation or an out-of-sample test of booking-policy profit.

Taken together, the numerical evidence supports the paper's structural claims at three levels. First, the direction-reversal and phase-transition experiments check the local economic mechanism. Second, the joint ambiguity grid shows that the two uncertainty layers affect policy through distinct channels and that their combination can make the service-risk constraint binding. Third, the finite-model audits and scaling experiment show that the proposed solution method produces reproducible, globally certified solutions for the discretized model. The matched monolithic benchmark strengthens this computational evidence by reproducing the same certified optimum to numerical precision while requiring far fewer instantiated constraints under delayed generation. Figure 8 makes the two complementary effects visible together: delayed generation compresses the finite robust model without changing the certified solution, while calibration uncertainty changes the booking mix and joint protection produces the largest portfolio displacement among the four canonical treatments. Finally, the public-data illustration grounds the segment-level cancellation and load primitives in observable booking records while deliberately stopping short of implemented policy-performance evidence. The experiments therefore do not establish realized profit gains for any particular hotel.

8. Discussion

8.1. Theoretical implications

The central theoretical implication is that calibration uncertainty has an endogenous adverse direction. A uniform conservative shift in all show-up rates is generally not the correct robust response. The adverse direction depends on the comparison between a segment's incremental service value and its marginal capacity burden. This comparison creates a phase transition: The same segment may be protected against downward show-up miscalibration in one operating regime and against upward show-up miscalibration in another. The bidirectional rule therefore converts heterogeneous cancellation risk into a decision-dependent optimization object rather than treating it as a fixed predictive input.

The extreme-point and budget allocation results clarify how this uncertainty enters the robust counterpart and the solution algorithm. The adversary concentrates its budget on accepted segments with wide calibration intervals and large economic sensitivity. This extends the price of robustness logic of budgeted uncertainty [15] by allowing the adverse sign to depend endogenously on segment value, accepted quantity, and marginal capacity pressure. The box-local expression provides an interpretable coordinatewise diagnostic, whereas the active scenario stationarity condition provides the continuous relaxation certificate under weighted budgeted ambiguity. The reported experiments use that continuous relaxation. If integer booking limits are imposed, the same finite counterpart becomes an MILP and certification is then based on the corresponding finite-model optimality gap; integer versus continuous performance is not claimed as an empirical result here.

The computational evidence also clarifies why the structural propositions matter. Direction reversal and the phase transition explain the sign and operating regime dependence of adverse calibration error; the extreme-point and weighted budget results make exact separation possible; and the matched benchmark shows that this reduction reproduces the monolithic finite-model optimum while instantiating only 9–11 active cuts on the tested matched instances. The propositions therefore serve jointly as policy interpretation and as a computational reduction, rather than as isolated formal statements.

8.2. Hotel policy interpretation

In the hotel cost-structured application, the robust model should not be read as a rule that simply suppresses total booking volume. The numerical results show that robustness is obtained mainly through segment-selective protection and worst-case load evaluation. Vacancy cost, overflow cost, cancellation-retention value, and the robust CVaR service constraint jointly determine whether a segment should be protected, reduced, or retained. This explains why a high-value segment can be harmed by downward show-up error in a low-pressure regime but by upward show-up error in a congested regime. The same reservation class can therefore change its calibration risk regime as occupancy pressure or service recovery cost changes.

The calibration budget Γ and Wasserstein radius ε also have different managerial meanings. The calibration budget represents confidence in segment-level show-up probabilities and primarily changes the portfolio's exposure to segment-level miscalibration. The public data illustration supports this interpretation only at the level of calibration range grounding: It shows substantial differences in segment cancellation and show-up rates, but it does not provide implemented booking policy evidence. The Wasserstein radius represents confidence in the conditional aggregate load law and primarily changes tail load protection, consistent with variation regularization and transport duality interpretations of Wasserstein DRO [21–23]. The representative policy experiment makes this distinction concrete. Calibration protection alone costs 189.997 synthetic units and changes the expected load, whereas distributional protection alone imposes the exact fixed-scale Wasserstein objective charge $\varepsilon\sigma c_o = 336$ and the exact CVaR charge $\varepsilon\sigma/(1-\beta) = 8$ at $\varepsilon = 0.1$, leaving the nominal acceptance portfolio unchanged in this fixed-scale specialization. That portfolio invariance is not a generic property of Wasserstein DRO when the residual scale depends on the decision. Under joint ambiguity, the protection cost is 546.269 units, the booking mix changes, and the robust CVaR limit binds. The calibration component conditional on $\varepsilon = 0.1$ is 210.269 units, compared with 189.997 at $\varepsilon = 0$, illustrating an interaction created by reoptimization under the joint uncertainty set.

8.3. Limitations

Several limitations should be emphasized. First, the main experiments are controlled cost-structured experiments. They check the structural claims, uncertainty set sensitivity, and finite optimization procedure, but they do not estimate causal effects or realized hotel profits. The public booking records are used only as a real-world-calibrated illustration of segment shares, cancellation levels, revenue proxies, stay load proxies, and split-sample calibration bands. They do not contain true capacity decisions, rejected demand, walking costs, or realized overbooking control actions, and therefore cannot support claims about implemented hotel policy performance.

Second, the model is static. It omits sequential arrivals, endogenous prices, customer choice, upgrades, dynamic cancellations, and multiday network recourse. These features are important in full-scale hotel revenue management, but adding them would obscure the paper's main objective: To isolate the structural effect of bidirectional calibration uncertainty under asymmetric capacity loss. Third, the truncated skew-normal law is a reference family for the standardized aggregate load residual, not a universal claim about all hotel load processes. Finally, the finite convergence result relies on exact solution of the finite master and adversarial subproblems, a polyhedral calibration set, and the chosen quadrature and piecewise-linear resolution. The scalability experiment covers 8–96 segments at a fixed numerical resolution; it supports stable performance over that tested range but does not establish an asymptotic complexity rate. The controlled policy experiments use the continuous relaxation, so fractional x_j values should be interpreted as segment-level booking control allocation levels rather than literal fractions of individual reservations; implementation with integer booking limits requires the MILP version of the finite counterpart. The joint out-of-sample statement in Theorem 5.9 is also conditional on the statistical coverage of both ambiguity sets: this study does not estimate η or ζ, nor does it select Γ and ε from a formal deployment-level coverage procedure. Finally, the matched monolithic benchmark deliberately uses 3001 reference quantiles and a 1001-point load grid to keep the full counterpart tractable. Its objective values should therefore be used only for the formulation-level comparison, not mixed numerically with the higher-resolution policy grid.

9. Conclusion

This paper develops a bidirectional calibration-robust booking control framework for reservation-based capacity systems. Segment-level show-up probabilities are allowed to vary within a weighted budgeted calibration set, while aggregate residual load is protected through a Wasserstein neighborhood of a standardized truncated skew-normal reference law. The decision maker chooses a heterogeneous acceptance portfolio under asymmetric vacancy and overflow costs and a robust CVaR service-risk constraint.

The main methodological contribution is to show that heterogeneous cancellation risk matters through the optimization value of probability error, not merely through predictive heterogeneity. The analysis establishes a calibration direction reversal, a segment phase transition, an extreme-point adversary, a local adversarial budget allocation rule, a closed-form box-local diagnostic, an active scenario KKT certificate for weighted budgeted ambiguity, and a bound on the price of calibration uncertainty. It also provides convexity and critical fractile results, normalized policy equivalence, and a finite convergence scenario and piece cutting-plane algorithm for the finite piecewise-linear model.

The computational study supports these claims in a controlled setting. The uncertainty-grid results quantify the separate and joint effects of calibration protection and Wasserstein ambiguity. The compact algorithmic efficiency summary shows that all tested scalability runs are solved with stable iteration counts, small cut counts, sub-second solution times, and zero reported residuals at the given numerical precision. The matched benchmark further reproduces the monolithic finite-model optimum within numerical tolerance while using only 9–11 generated cuts instead of 18,009–66,033 full constraints over $J \in \{8, 16, 32\}$; Figure 8 summarizes this computational compression together with the uncertainty induced objective and portfolio responses. A real-world-calibrated illustration based on public hotel booking records further grounds the segment shares, cancellation levels, show-up rates, revenue proxies, stay load proxies, and calibration bands used to interpret the model. This public data component is intentionally limited to calibration range grounding and should not be read as evidence of implemented hotel policy performance.

Overall, the paper contributes a structural and computational optimization theory for calibration-aware booking control. Its practical value is to separate three objects that are often conflated in reservation analytics: Predictive show-up heterogeneity, segment-level calibration uncertainty, and aggregate load distributional ambiguity. This separation yields interpretable protection rules, auditable active scenario certificates, and a measurable price of calibration uncertainty for cost-structured booking decisions.

Data availability

The controlled numerical instances used in the main computational study are synthetic and were generated by the authors to evaluate the structural and algorithmic properties of the proposed robust booking control model. The public data illustration uses a publicly available hotel booking data set hosted on Kaggle [48], which is described by its provider as an updated and cleaned version of the Hotel Booking Demand data set originally documented by Antonio et al. [49]. The public records are used only to construct segment-level calibration summaries and are not used as field evidence of implemented booking control performance.

Author contributions

Conceptualization, methodology, mathematical modeling, formal analysis, software design, and writing—original draft, J.C.; reservation management interpretation, synthetic policy design, computational verification, and writing—review and editing, S.Y.L. All authors have read and agreed to the submitted version of the manuscript.

Use of generative-AI tools declaration

The authors declare that generative-artificial intelligence (AI) tools were used only to assist with language editing, formatting refinement, and drafting support during manuscript preparation. All modeling choices, mathematical derivations, computational experiments, interpretation, and final scientific judgments were performed and verified by the authors. The authors take full responsibility for the content of this article.

Acknowledgments

This research was supported by a Research Grant from Kyungwoon University in 2026. The authors gratefully acknowledge this support.

Conflict of interest

The authors declare no conflicts of interest.

References

1. Y. Feng, B. Xiao, A continuous-time seat control model for single-leg flights with no-shows and optimal overbooking upper bound, *Eur. J. Oper. Res.*, **174** (2006), 1298–1316. https://doi.org/10.1016/j.ejor.2005.05.008
2. A. K. Strauss, R. Klein, C. Steinhardt, A review of choice-based revenue management: Theory and methods, *Eur. J. Oper. Res.*, **271** (2018), 375–387. https://doi.org/10.1016/j.ejor.2018.01.011
3. K. Lai, W. Ng, A stochastic approach to hotel revenue optimization, *Comput. Oper. Res.*, **32** (2005), 1059–1072. https://doi.org/10.1016/j.cor.2003.09.012
4. S. Liu, K. K. Lai, S. Y. Wang, Booking models for hotel revenue management considering multiple-day stays, *Int. J. Revenue Manag.*, **2** (2008), 78–91. https://doi.org/10.1504/IJRM.2008.018179
5. N. Aydın, S. I. Birbil, J. B. G. Frenk, N. Noyan, Single-leg airline revenue management with overbooking, *Transp. Sci.*, **47** (2013), 560–583. https://doi.org/10.1287/trsc.1120.0444
6. D. D. Sierag, G. M. Koole, R. D. V. D. Mei, J. I. V. D. Rest, B. Zwart, Revenue management under customer choice behaviour with cancellations and overbooking, *Eur. J. Oper. Res.*, **246** (2015), 170–185. https://doi.org/10.1016/j.ejor.2015.04.014
7. N. Antonio, A. D. Almeida, L. Nunes, Predicting hotel booking cancellations to decrease uncertainty and increase revenue, *Tour. Manag. Stud.*, **13** (2017), 25–39. https://doi.org/10.18089/tms.2017.13203
8. N. Antonio, A. D. Almeida, L. Nunes, Big data in hotel revenue management: Exploring cancellation drivers to gain insights into booking cancellation behavior, *Cornell Hosp. Q.*, **60** (2019), 298–319. https://doi.org/10.1177/1938965519851466
9. S. Chen, E. W. T. Ngai, Y. Ku, Z. Xu, X. Gou, C. Zhang, Prediction of hotel booking cancellations: Integration of machine learning and probability model based on interpretable feature interaction, *Decis. Support Syst.*, **170** (2023), 113959. https://doi.org/10.1016/j.dss.2023.113959

10. Z. Liu, K. W. D. Bock, L. Zhang, Explainable profit-driven hotel booking cancellation prediction based on heterogeneous stacking-based ensemble classification, *Eur. J. Oper. Res.*, **321** (2025), 284–301. https://doi.org/10.1016/j.ejor.2024.08.026

11. A. Ampountolas, M. Legg, A comprehensive approach to enhancing short-term hotel cancellation forecasts through dynamic machine learning models, *Tour. Econ.*, **32** (2026), 321–341. https://doi.org/10.1177/13548166251318768

12. V. Pimentel, A. Aizezikali, T. Baker, Hotel revenue management: Benefits of simultaneous overbooking and allocation problem formulation in price optimization, *Comput. Ind. Eng.*, **137** (2019), 106073. https://doi.org/10.1016/j.cie.2019.106073

13. Q. Zhai, Y. Tian, J. Luo, J. Zhou, Hotel overbooking based on no-show probability forecasts, *Comput. Ind. Eng.*, **180** (2023), 109226. https://doi.org/10.1016/j.cie.2023.109226

14. A. V. D. Boer, D. D. Sierag, Decision-based model selection, *Eur. J. Oper. Res.*, **290** (2021), 671–686. https://doi.org/10.1016/j.ejor.2020.08.025

15. D. Bertsimas, M. Sim, The price of robustness, *Oper. Res.*, **52** (2004), 35–53. https://doi.org/10.1287/opre.1030.0065

16. E. Delage, Y. Ye, Distributionally robust optimization under moment uncertainty with application to data-driven problems, *Oper. Res.*, **58** (2010), 595–612. https://doi.org/10.1287/opre.1090.0741

17. P. M. Esfahani, D. Kuhn, Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations, *Math. Program.*, **171** (2018), 115–166. https://doi.org/10.1007/s10107-017-1172-1

18. D. Kuhn, P. Mohajerin Esfahani, V. A. Nguyen, S. Shafieezadeh-Abadeh, Wasserstein distributionally robust optimization: Theory and applications in machine learning, *Oper. Res. Manag. Sci. Age Anal.*, 2019. https://doi.org/10.1287/educ.2019.0198

19. G. A. Hanasusanto, D. Kuhn, Conic programming reformulations of two-stage distributionally robust linear programs over Wasserstein balls, *Oper. Res.*, **66** (2018), 849–869. https://doi.org/10.1287/opre.2017.1698

20. R. Gao, A. J. Kleywegt, Distributionally robust stochastic optimization with wasserstein distance, *Math. Oper. Res.*, **48** (2023), 603–655. https://doi.org/10.1287/moor.2022.1275

21. R. Gao, X. Chen, A. J. Kleywegt, Wasserstein distributionally robust optimization and variation regularization, *Oper. Res.*, **72** (2024), 1177–1191. https://doi.org/10.1287/opre.2022.2383

22. L. Zhang, J. Yang, R. Gao, A short and general duality proof for wasserstein distributionally robust optimization, *Oper. Res.*, **73** (2025), 2146–2155. https://doi.org/10.1287/opre.2023.0135

23. J. Zhen, D. Kuhn, W. Wiesemann, A unified theory of robust and distributionally robust optimization via the primal-worst-equals-dual-best principle, *Oper. Res.*, **73** (2025), 862–878. https://doi.org/10.1287/opre.2021.0268

24. Y. Fan, H. Park, G. Xu, Quantifying distributional model risk in marginal problems via optimal transport, *Math. Oper. Res.*, (2025). https://doi.org/10.1287/moor.2024.0557

25. A. N. Elmachtoub, P. Grigas, Smart "Predict, then optimize", *Manag. Sci.*, **68** (2022), 9–26. https://doi.org/10.1287/mnsc.2020.3922

26. M. Sim, Q. Tang, M. Zhou, T. Zhu, The analytics of robust satisficing: predict, optimize, satisfice, then fortify, *Oper. Res.*, **73** (2025), 2708–2728. https://doi.org/10.1287/opre.2023.0199

27. R. T. Rockafellar, S. Uryasev, Optimization of conditional value-at-risk, *J. Risk*, **2** (2000), 21–41. https://doi.org/10.21314/JOR.2000.038

28. Y. Lan, M. O. Ball, I. Z. Karaesmen, J. X. Zhang, G. X. Liu, Analysis of seat allocation and overbooking decisions with hybrid information, *Eur. J. Oper. Res.*, **240** (2015), 493–504. https://doi.org/10.1016/j.ejor.2014.07.021

29. P. Goldman, R. Freling, K. Pak, N. Piersma, Models and techniques for hotel revenue management using a rolling horizon, *J. Revenue Pricing Manag.*, **1** (2002), 207–219. https://doi.org/10.1057/palgrave.rpm.5170025

30. B. Vinod, Unlocking the value of revenue management in the hotel industry, *J. Revenue Pricing Manag.*, **3** (2004), 178–190. https://doi.org/10.1057/palgrave.rpm.5170105

31. S. Ivanov, V. Zhechev, Hotel revenue management: A critical literature Review, *Tourism*, **60** (2012), 175–197. https://doi.org/0.2139/ssrn.1977467

32. C. C. Chen, Z. Schwartz, P. Vargas, The search for the best deal: How hotel cancellation policies affect the search and booking decisions of deal-seeking customers, *Int. J. Hosp. Manag.*, **30** (2011), 129–135. https://doi.org/10.1016/j.ijhm.2010.03.010

33. Z. Schwartz, T. D. Webb, M. Altin, A. Riasi, Overbooking and performance in hotel revenue management, *Int. J. Hosp. Manag.*, **129** (2025), 104192. https://doi.org/10.1016/j.ijhm.2025.104192

34. N. Antonio, A. D. Almeida, L. Nunes, Predicting hotel bookings cancellation with a machine learning classification model, *Proc. Int. Conf. Mach. Learn. Appl.*, (2017), 1049–1054. https://doi.org/10.1109/ICMLA.2017.00-11

35. A. J. Sánchez-Medina, E. C-Sánchez, Using machine learning and big data for efficient forecasting of hotel booking cancellations, *Int. J. Hosp. Manag.*, **89** (2020), 102546. https://doi.org/10.1016/j.ijhm.2020.102546

36. M. Yoo, A. K. Singh, N. Loewy, Predicting hotel booking cancelation with machine learning techniques, *J. Hosp. Tour. Technol.*, **15** (2024), 54–69. https://doi.org/10.1108/JHTT-07-2022-0227

37. Z. Liu, P. Jiang, J. Wang, Z. Du, X. Niu, L. Zhang, Hospitality order cancellation prediction from a profit-driven perspective, *Int. J. Contemp. Hosp. Manag.*, **35** (2023), 2084–2112. https://doi.org/10.1108/IJCHM-06-2022-0737

38. P. C. Silvestre, N. António, P. Carrasco, Navigating uncertainty: Enhancing hotel cancellation predictions with adaptive machine learning, *Inf. Technol. Tour.*, **28** (2026), 9. https://doi.org/10.1007/s40558-025-00349-9

39. J. Gönsch, A survey on risk-averse and robust revenue management, *Eur. J. Oper. Res.*, **263** (2017), 337–348. https://doi.org/10.1016/j.ejor.2017.05.033

40. F. Lin, X. Fang, Z. Gao, Distributionally robust optimization: A review on theory and applications, *Numer. Algebra Control Optim.*, **12** (2022), 159–212. https://doi.org/10.3934/naco.2021057

41. K. S. Shehadeh, R. Padman, A distributionally robust optimization approach for stochastic elective surgery scheduling with limited intensive care unit capacity, *Eur. J. Oper. Res.*, **290** (2021), 901–913. https://doi.org/10.1016/j.ejor.2020.09.001

42. H. Takagi, Three-class static revenue management model with buy-up behavior, *J. Ind. Manag. Optim.*, **20** (2024), 548–569. https://doi.org/10.3934/jimo.2023090

43. A. Azzalini, A class of distributions which includes the normal ones, *Scand. J. Stat.*, **12** (1985), 171–178. https://doi.org/10.6092/ISSN.1973-2201/711

44. M. Shaked, J. G. Shanthikumar, *Stochastic orders*, Springer, New York, 2007. https://doi.org/10.1007/978-0-387-34675-5

45. J. F. Bonnans, A. Shapiro, *Perturbation analysis of optimization problems*, Springer, New York, 2000. https://doi.org/10.1007/978-1-4612-1394-9

46. R. T. Rockafellar, R. J. Wets, *Variational analysis*, Springer, Berlin, (1998). https://doi.org/10.1007/978-3-642-02431-3

47. C. Peng, E. Delage, Data-driven optimization with distributionally robust second-order stochastic dominance constraints, *Oper. Res.*, **72** (2024), 1298–1316. https://doi.org/10.1287/opre.2022.2387https://doi.org/10.1287/opre.2022.2387

48. K. S. Bedmutha, Hotel booking reservation, 2024. Available from: https://www.kaggle.com/datasets/kundanbedmutha/hotel-booking-reservation.

49. N. Antonio, A. D. Almeida, L. Nunes, Hotel booking demand datasets, *Data Brief*, **22** (2019), 41–49. https://doi.org/10.1016/j.dib.2018.11.126