UNIVERSITAT D'ALACANT
UNIVERSIDAD DE ALICANTE
Facultat de Ciències Econòmiques i Empresarials
Facultad de Ciencias Económicas y Empresariales

**MASTER IN ECONOMICS WITH DATA SCIENCE**

**ACADEMIC YEAR 2025 - 2026**

# DYNAMIC PRICING UNDER CAPACITY CONSTRAINTS: AN AGENT-BASED APPROACH TO COMPETITION IN THE CAR RENTAL INDUSTRY

AUTHOR: DANIEL SERRANO ALZAMORA

TUTOR: ANDREEA CORINA DOBRE
REVENUE & PRICING DEPARTMENT

San Vicente del Raspeig, September 2026

# Abstract

This work uses an Agent-Based Model with synthetic data to evaluate whether a pricing rule based on the relative competitive position against a small subset of price-adjacent rivals (Top-3) outperforms rules based on aggregate market indicators or on own occupancy, in terms of revenue and occupancy stability, within a capacity-constrained oligopolistic car rental market. The motivation stems from an observation, during an internship in the Pricing and Revenue department of a car rental company, that this heuristic appeared operationally more useful than tracking market aggregates. A ten-firm market is simulated under four pricing rules and four combinations of capacity tension and fleet size dispersion. Results show that the positional strategy and the market-average strategy show no statistically significant differences in RevPAR across any of the four scenarios, although no strategy shows a systematic advantage in occupancy stability. Average RevPAR consistently doubles, across strategies, from low to high season. The most robust finding is that a hybrid strategy combining the positional signal with own occupancy fails to outperform either of the two best-performing pure strategies ($p < 0.0001$ across all comparisons), suggesting that, under the adopted weighting and calibration, incorporating the occupancy signal reduces the performance of the positional rule. Model limitations, synthetic calibration, and implications for revenue management practice in the sector are discussed.


**Keywords:** dynamic pricing, agent-based models, revenue management, capacity-constrained competition, car rental.

# Resumen

Este trabajo evalúa, mediante un Modelo Basado en Agentes con datos sintéticos, si una regla de *pricing* basada en la posición competitiva relativa frente a un subconjunto reducido de competidores próximos en precio (Top-3) supera, en ingreso y estabilidad de ocupación, a reglas basadas en indicadores agregados de mercado o en la ocupación propia, en un mercado oligopolístico de alquiler de vehículos con restricción de capacidad. La motivación surge de la observación, durante un periodo de prácticas en el departamento de *Pricing y Revenue* de una empresa de alquiler de vehículos, de que dicha heurística parecía más útil operativamente que el seguimiento de agregados de mercado. Se simula un mercado de diez empresas bajo cuatro reglas de *pricing* y cuatro combinaciones de tensión de capacidad y tamaño de flota. Los resultados muestran que la estrategia posicional y la basada en la media de mercado no muestran diferencias estadísticamente significativas en RevPAR en ninguno de los cuatro escenarios, si bien no se observa una ventaja sistemática de ninguna estrategia en términos de estabilidad de ocupación. El RevPAR medio se duplica, de forma consistente entre estrategias, al pasar de temporada baja a temporada alta. El hallazgo más robusto del trabajo es que una estrategia híbrida, que combina la señal posicional con la ocupación propia, no supera a ninguna de las dos estrategias de mejor desempeño por separado ($p < 0.0001$ en todas las comparaciones), lo que sugiere que, bajo la ponderación y calibración adoptadas, incorporar la señal de ocupación reduce el rendimiento de la regla posicional. Se discuten las limitaciones del modelo, su calibración sintética y las implicaciones para la práctica de *revenue management* en el sector.

**Palabras clave:** pricing dinámico, modelos basados en agentes, revenue management, competencia con restricción de capacidad, alquiler de vehículos.

# Contents

# 1 Introduction

The car rental industry (*rent-a-car*) exhibits a set of characteristics that make it a particularly suitable case study for analyzing dynamic pricing under capacity constraints: inventory (the fleet) is fixed in the short term, perishable in the sense that an unrented vehicle on a given day represents irrecoverable revenue (Talluri and van Ryzin, 2004 [1]), and competition among operators takes place primarily through published prices, which are visible and comparable across firms in real time. These conditions have traditionally motivated the application of *revenue management* techniques in related industries—such as airlines and hotels—yet present unique features in car rental: rental duration is shorter and more variable, the number of direct competitors in a local market is small and observable, and pricing decisions are typically made on a daily basis, often decentralized among pricing analysts who combine heuristic rules with subjective judgment, rather than through the explicit resolution of a dynamic optimization problem.

This study arises from observing this operational practice during an internship in the *Pricing and Revenue* department of a car rental company. In that setting, it was striking that, for daily pricing decisions, visual comparison against a small subset of price-adjacent competitors—termed in this work the *Top-3 adjacent competitors*—appeared operationally to provide a more immediate and actionable signal than tracking aggregate market statistics, such as the competitive block's average price. This empirical observation, qualitative and anecdotal in origin, raises a question amenable to formal analysis: **Under what market conditions does a pricing rule based on relative competitive position against a subset of adjacent rivals outperform, in economic performance, rules based on aggregate market indicators or own fleet occupancy?**

For the purposes of the experimental design (Section 3.6), this general inquiry breaks down into three specific research questions: whether the Positional strategy generates a higher RevPAR than the market-average strategy (RQ1); whether it improves occupancy stability relative to the own-occupancy strategy (RQ2); and whether combining both signals improves performance relative to the constituent pure rules (RQ3). These three questions are addressed in Sections 4.2, 4.3, and 4.4, respectively. A fourth, exploratory finding outside the original experimental design—whether the advantage of positional information depends on demand price sensitivity, tested via a systematic parameter sweep (Section 4.2)—emerged during the study and is discussed in Section 4.2 as a secondary result.

Due to confidentiality restrictions, this work cannot use proprietary data or reproduce proprietary pricing algorithms. Instead, it develops an Agent-Based Model (ABM) with synthetic data representing a stylized oligopolistic car rental market where virtual firms compete on price under fleet capacity constraints following distinct decision rules. The aim is not to reproduce numerically the behavior of any real-world market, but to isolate and compare, in a controlled environment, the underlying economic mechanism: whether monitoring price-adjacent competitors—compared to tracking the market as a whole or relying solely on own occupancy—systematically yields superior outcomes in terms of revenue and inventory stability, and under what capacity and seasonal demand conditions this advantage persists, diminishes, or reverses.

Choosing an ABM over analytical approaches—typically, capacity-constrained Bertrand-Edgeworth models (Kreps and Scheinkman, 1983 [2]; Davidson and Deneckere, 1986 [3])—is motivated by the fact that analytical models rarely admit pure-strategy equilibria when agent decision rules are heterogeneous, and require full-rationality assumptions inconsistent with the bounded, heuristic nature of observed pricing decisions (Simon, 1955 [4]). Compared to empirical econometrics on real data, an ABM allows controlled *counterfactual* comparisons across decision rules—evaluating what would have occurred had a firm adopted an alternate rule, holding all else constant—which cannot be identified from observational data generated by a single historical policy. This methodological rationale is detailed in Section 2.

The remainder of this thesis is structured as follows. Section 2 reviews the theoretical framework across three pillars: price competition under capacity constraints, revenue management and dynamic pricing, and agent-based market models. Section 3 formalizes the simulation model: agents, demand function, the four evaluated pricing rules, and explicit modeling assumptions. Section 4 presents and analyzes the simulation results across capacity and seasonality scenarios. Section 5 summarizes findings, discusses limitations, and outlines practical implications for revenue management. Finally, Section 6 proposes six potential avenues for future research.

# 2 Theoretical Framework

This section outlines the theoretical foundations across three areas supporting the methodological approach: price competition under capacity constraints, revenue management applied to the car rental industry, and agent-based computational economics. Each subsection concludes by highlighting its specific contribution to the model design formalized in Section 3.

## 2.1 Price Competition Under Capacity Constraints

The natural starting point for analyzing competition among a small number of car rental operators is the Bertrand pricing model, where firms compete via prices and demand is entirely captured by the lowest-price operator. However, Bertrand's implicit premise—that any firm can fulfill all demand it faces—breaks down in car rental, where fleet capacity is fixed in the short term. Under explicit capacity limits, the competitive Bertrand outcome ($p = mc$) fails as an equilibrium: if all firms set marginal cost prices, the cheapest operator exhausts capacity immediately, leaving residual demand for remaining firms and providing incentives to raise prices. Kreps and Scheinkman (1983 [2]) demonstrate that under specific rationing assumptions for unfulfilled demand, a two-stage game—capacity selection followed by price competition—converges to the Cournot outcome rather than Bertrand. Davidson and Deneckere (1986 [3]) qualify this finding by showing its sensitivity to rationing rules: under proportional rationing—the convention adopted in this study, in contrast to "efficient" consumer-surplus rationing—the equilibrium moves further from the competitive benchmark, enhancing capacity-derived market power.

This literature justifies, first, why capacity constraints represent a core economic mechanism rather than a secondary implementation detail: without capacity limits, matching the cheapest competitor triggers a price war collapsing to marginal cost. Capacity ceilings create the operational space where distinct heuristics—tracking market averages, own occupancy, or adjacent rivals—generate differentiated revenue outcomes. Second, it justifies the simulation approach: as discussed in Section 2.3, introducing heterogeneous decision rules—none of which are classical profit maximizers across our four specifications—precludes analytical closed-form solutions beyond the symmetric, fully rational benchmarks in classical literature.

## 2.2 Revenue Management and Dynamic Pricing in Car Rental

Applying revenue management to car rental dates back to the early 1990s. Carroll and Grimes (1995 [6]) describe implementing yield management systems at Hertz, adapting established airline practices to fleet operations. The most prominent case in the literature is National Car Rental: Geraghty and Johnson (1997 [5]) document how a comprehensive revenue management system combining capacity controls, pricing, and reservation management helped reverse the company's financial crisis in the mid-1990s, driving measurable revenue gains. This case is relevant here because it establishes that occupancy-driven dynamic pricing (corresponding to our *Occupancy* rule) is the historically dominant industry benchmark. In contrast, competitive-position rules—closer to the operational practices observed during the internship motivating this thesis—remain less explored in academic literature despite widespread operational use.

The overarching conceptual framework of revenue management—managing fixed, perishable inventory under demand uncertainty via price and availability controls—is synthesized by Talluri and van Ryzin (2004 [1]). Characterizing the problem as the optimal control of *fixed capacity, perishable inventory, and segmentable demand* applies directly to car rental: an unrented vehicle on a given day represents an irrecoverable loss, directly analogous to an unsold airline seat or hotel room. Subsequent car rental literature has focused predominantly on network fleet allocation and stochastic reservation control models. While sharing revenue maximization objectives, these address fleet distribution and availability rather than daily published price adjustments against competitors.

This domain contributes two main elements: first, it establishes car rental pricing within perishable-inventory revenue management, justifying RevPAR as the primary metric, parallel to lodging RevPAR or airline yield; second, it reveals that existing research focuses on centralized optimization and occupancy-based rules rather than relative positioning against adjacent rivals, highlighting the relevance of evaluating localized competitive tracking.

## 2.3 Agent-Based Computational Models Applied to Markets

Agent-Based Computational Economics (ACE) is defined, following Tesfatsion (2006 [7]), as the computational study of economic processes modeled as dynamic systems of interacting agents, where aggregate regularities emerge from repeated local interactions rather than analytical equilibrium imposed *ex ante*. This perspective aligns with recent literature on algorithmic pricing in competitive environments, which explores how simple heuristic pricing rules generate non-trivial market dynamics under repeated interaction (Calvano et al., 2020 [12]). This approach is well-suited for settings with asymmetric information, imperfect competition, and multiple potential equilibria that resist tractable closed-form solutions, matching the conditions outlined in Section 2.1.

A core element of ACE is the representation of boundedly rational agents (Simon, 1955 [4]) operating through heuristics rather than solving global dynamic optimization problems. Heterogeneous Agent Models in financial economics—reviewed by LeBaron (2006 [8])—show that interactions among simple, diverse decision rules generate aggregate market phenomena, such as oscillations and regime-dependent rule advantages, absent in representative-agent frameworks. This structure maps directly to our design: the four pricing rules (Section 3.3) represent boundedly rational heuristics utilizing different information subsets. We evaluate them counterfactually across symmetric markets—where all market operators adopt the same rule—to analyze emerging aggregate dynamics in RevPAR and occupancy stability across varying capacity and seasonality regimes.

This framework completes the methodological justification: unlike analytical approaches (Section 2.1), which require full-rationality assumptions, and empirical econometrics, constrained by proprietary data access and lack of exogenous rule variation, an ABM allows controlled counterfactual analysis across identical synthetic market environments.

# 3 Simulation Model Design

## 3.1 Notation and Temporal Framework

The Agent-Based Model is implemented in Python[1]. We consider an oligopolistic market consisting of $N = 10$ virtual firms $i \in \{1, \dots, N\}$ competing on price across a discrete time horizon $t \in \{1, \dots, T\}$ (days). Each firm $i$ operates a fixed fleet of capacity $K_i$ (number of vehicles), held constant throughout the simulation. Setting $N = 10$ rather than a smaller value ensures that the "Top-3 adjacent competitors" defining the *Positional* strategy (Section 3.3) does not overlap substantially with the entire market. With $N = 10$, filtering the Top-3 excludes 6 of 9 rivals for each firm, preserving a meaningful distinction between localized and aggregate market tracking.

The state of each firm at step $t$ is defined by:

- $p_{i,t}$: daily rental price set by firm $i$.
- $O_{i,t} \in [0, K_i]$: occupied (rented) vehicles at the start of day $t$.
- $o_{i,t} = O_{i,t}/K_i \in [0, 1]$: fleet occupancy rate.

Firm heterogeneity in the baseline model is limited to fleet capacity $K_i$, drawn synthetically from a bounded discrete uniform distribution ($K_i \sim \mathcal{U}\{80, 150\}$ vehicles), generating

[1] The complete replication code and Jupyter notebooks are publicly available at `https://github.com/Daniseralz/TFM-Dynamic-Pricing-ABM`.

varied operator sizes without imposing structural dominance by construction.

The sequence of events in each period $t$ proceeds as follows:

1. Each firm observes market states at $t-1$ (prices and occupancy across all firms) and sets price $p_{i,t}$ according to its pricing rule (Section 3.3).

2. Aggregate market demand $D_t$ is realized and allocated across firms via a logit choice model (Section 3.2).

3. Proportional capacity rationing is applied to excess demand where applicable.

4. Fleet occupancy $O_{i,t+1}$ updates by adding new bookings and subtracting vehicle returns from rentals initiated $L$ days earlier.

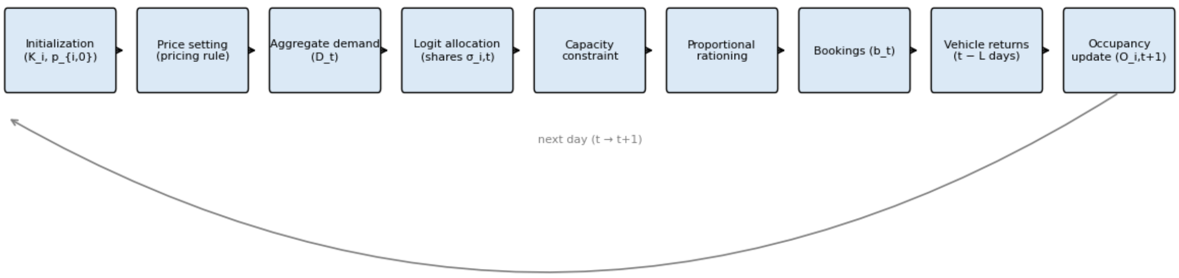


**Figure 1:** *Sequence of events within the simulation model for each period $t$.*

Figure 1 illustrates the procedural sequence described above.

## 3.2 Demand Function and Rationing Mechanism

Aggregate daily market demand at day $t$ is specified as

$$D_t = D_0 \cdot s_t \cdot (1 + \varepsilon_t), \qquad \varepsilon_t \sim \mathcal{N}(0, \sigma_D^2) \tag{1}$$

where $D_0$ is baseline daily demand, $s_t \in \{s_{\text{low}}, s_{\text{high}}\}$ is a seasonal multiplier ($s_{\text{low}} = 1.0$ and $s_{\text{high}} = 1.6$), and $\varepsilon_t$ represents zero-mean Gaussian noise capturing price-independent demand shocks ($\sigma_D = 0.05$).

Allocation of aggregate demand across firms follows a standard price *logit* model:

$$\sigma_{i,t} = \frac{\exp(-\beta\, p_{i,t})}{\sum_{j=1}^{N} \exp(-\beta\, p_{j,t})}, \qquad d_{i,t} = D_t \cdot \sigma_{i,t} \tag{2}$$

where $\beta > 0$ represents the price-sensitivity parameter (governing cross-price elasticity). This specification follows discrete choice random utility theory (McFadden, 1974 [10]; Anderson, de Palma, and Thisse, 1992 [11]). The parameter $\beta$ is calibrated in simulation so that realistic price differentials (10–15% relative to the market mean) yield substantial but stable market share variations.

Similarly, $D_0$ is calibrated against aggregate market capacity to produce a baseline natural occupancy around 55% during low season (documented in simulation code and discussed in Section 4.1).

As with standard multinomial logit models, this formulation satisfies the Independence of Irrelevant Alternatives (*IIA*) property: the ratio of choice probabilities between any two firms depends exclusively on their own prices. This simplification aligns with assuming zero brand or product quality differentiation across firms (Assumption 1, Section 3.5).

Actual demand *served* by firm $i$ is constrained by remaining fleet capacity:

$$b_{i,t} = \min\left(d_{i,t},\ K_i - O_{i,t}\right) \tag{3}$$

When one or more firms exhaust available capacity ($O_{i,t} = K_i$), unserved demand $d_{i,t} - b_{i,t}$ is redistributed **proportionally** among firms with remaining availability, weighted by their relative logit shares. This follows standard proportional rationing conventions in capacity-constrained Bertrand-Edgeworth models without priority ordering across consumer types (Davidson and Deneckere, 1986 [3]; Section 2.1).

Rented vehicles are returned after a fixed average rental duration $L$ (identical across all firms and bookings, with $L = 4$ days representing leisure rentals):

$$O_{i,t+1} = O_{i,t} + b_{i,t} - b_{i,t-L} \tag{4}$$

Daily revenue generated by firm $i$ at day $t$ is given by

$$R_{i,t} = b_{i,t} \cdot p_{i,t} \cdot L \tag{5}$$

where multiplying by $L$ reflects that each rental booked on day $t$ generates revenue across the entire rental contract duration (Equation 4). RevPAR (*Revenue per Available Car-Day*) is calculated over the steady-state period, excluding an initial BURN_IN = 20 days:

$$\text{RevPAR} = \frac{\sum_{t=\text{BURN\_IN}+1}^{T} \sum_{i=1}^{N} R_{i,t}}{\sum_{i=1}^{N} K_i \cdot T_{\text{effective}}}, \qquad T_{\text{effective}} = T - \text{BURN\_IN} \tag{6}$$

## 3.3 Pricing Rules

All pricing strategies share an identical bounded daily adjustment structure driven by a strategy-specific signal $\Delta_{i,t} \in [-1, 1]$ and a maximum daily adjustment rate $\lambda$ (set to $\lambda = 0.03$, or 3% daily):

$$p_{i,t+1} = \text{clip}\Big(p_{i,t} \cdot (1 + \lambda\, \Delta_{i,t}),\ p_{\min},\ p_{\max}\Big) \tag{7}$$

This shared formulation ensures direct comparability across strategies, isolating differences in information usage from adjustment step sizing.

Because simulations begin with an empty market ($O_{i,0} = 0$, Section 3.1), the first 20 days serve as a fleet-filling burn-in period and are excluded from metric calculations (Section 4.1), analyzing the remaining 160 steady-state days.

$$\textbf{Average:} \quad \Delta_{i,t}^{M} = \frac{\bar{p}_{-i,t} - p_{i,t}}{\bar{p}_{-i,t}} \qquad \bar{p}_{-i,t} = \frac{1}{N-1} \sum_{j \neq i} p_{j,t} \tag{8}$$

$$\textbf{Occupancy:} \quad \Delta_{i,t}^{O} = \frac{o_{i,t} - o^*}{o^*} \qquad o^* = 0.85 \text{ (target occupancy)} \tag{9}$$

$$\textbf{Positional:} \quad \Delta_{i,t}^{P} = \frac{\bar{p}_{i,t}^{\text{Top3}} - p_{i,t}}{\bar{p}_{i,t}^{\text{Top3}}} \qquad \bar{p}_{i,t}^{\text{Top3}} = \text{mean price of the 3 rivals closest in price to } p_{i,t} \tag{10}$$

$$\textbf{Hybrid:} \quad \Delta_{i,t}^{H} = \gamma\, \Delta_{i,t}^{O} + (1-\gamma)\, \Delta_{i,t}^{P} \qquad \gamma = 0.5 \tag{11}$$

Simulations evaluate symmetric markets where **all** $N = 10$ **firms implement the same pricing rule**, comparing performance *across* symmetric markets rather than examining heterogeneous rule interactions within a single market. This benchmark design assesses the aggregate market outcome if an entire industry adopts a specific heuristic.

The *Positional* rule (10) defines adjacency strictly by absolute price difference $|p_{j,t} - p_{i,t}|$, selecting the 3 closest competitors among the 9 available rivals regardless of fleet size.

The term $\Delta^{O}_{i,t}$ is normalized relative to target occupancy $o^*$: under zero occupancy ($o_{i,t} = 0$), the signal reaches its theoretical lower bound $\Delta^{O}_{i,t} = -1$, triggering the maximum downward adjustment bounded by $\lambda$ (Equation 7).

## 3.4 Parameter Summary

Table 1 compiles the baseline numerical parameter values used throughout the model.

| Parameter | Symbol | Value | Description |
|---|---|---|---|
| Firms | $N$ | 10 | Number of market operators |
| Base demand | $D_0$ | 161 | Vehicles/day (calibrated, Section 3.2) |
| Price sensitivity | $\beta$ | 0.08 | Logit sensitivity parameter (calibrated) |
| Demand noise | $\sigma_D$ | 0.05 | Relative daily shock standard deviation |
| Low season | $s_{\text{low}}$ | 1.0 | Seasonal demand multiplier |
| High season | $s_{\text{high}}$ | 1.6 | Seasonal demand multiplier |
| Rental duration | $L$ | 4 | Rental duration in days (deterministic) |
| Max daily adjustment | $\lambda$ | 0.03 | Maximum daily adjustment rate (3%) |
| Target occupancy | $o^*$ | 0.85 | Target fleet occupancy (85%) |
| Hybrid weight | $\gamma$ | 0.5 | 50% occupancy / 50% positional |
| Capacity range | $K_i$ | [80, 150] | Fleet capacity bounds per firm |
| Time horizon | $T$ | 180 | Simulated days per run |
| Burn-in period | – | 20 | Days excluded from metric calculations |
| Replications | $N_{\text{seeds}}$ | 20 | Random seeds per scenario |

**Table 1:** *Reference numerical parameters of the simulation model.*

## 3.5 Explicit Assumptions

1. Market demand allocation depends solely on price differentials via the logit choice model (Equation 2); quality, brand reputation, and ancillary services are not modeled.

2. Unmet demand rationing follows a **proportional** allocation rather than an efficient consumer-surplus rule.

3. Rental duration $L$ is fixed, deterministic, and identical across all transactions.

4. Firms are **myopic**: prices adjust reactively based on localized heuristic signals without solving dynamic optimization programs or anticipating rival reactions (Simon, 1955 [4]; Section 2.3).

5. Fleet capacities $K_i$ remain fixed throughout the simulation horizon, excluding intra-period fleet transfers or secondary investments.

6. Vehicle marginal operational costs are omitted; outcomes evaluate gross operational revenue (RevPAR) rather than net operating profits.

7. Each firm quotes a single market-wide daily rate, without tariff tiers or advance booking yield curves (Talluri and van Ryzin, 2004 [1]).

8. Competitor prices are publicly observable in real time, but rival decision rules and internal fleet occupancies remain private.

## 3.6 Experimental Scenarios

The experimental design employs a factorial combination across two structural axes:

- **Capacity tension**: low season ($s_t = 1.0$, substantial capacity slack) versus high season ($s_t = 1.6$, tight capacity utilization).

- **Fleet size dispersion**: a homogeneous market (fleets tightly distributed around mean size) versus a dispersed market (coexistence of large and small operators), holding aggregate market capacity constant.

Each of the four pricing rules (Section 3.3) is simulated across the four scenario combinations, replicated over $N_{\text{seeds}} = 20$ distinct random seed runs for demand shocks $\varepsilon_t$ (Equation 1), reporting mean metrics alongside across-seed dispersion.

## 3.7 Internal Model Verification

Because empirical calibration against proprietary market data was restricted (Section 3.8), we conducted consistency verification tests on simulator components to confirm expected properties of logit choice dynamics (Equation 2) and rationing mechanisms (Section 3.2):

1. Identical prices produce uniform market shares equal to $1/N$ across all operators.
2. Unilaterally increasing a firm's price, *ceteris paribus*, strictly decreases its market share.
3. As $\beta \to 0$ (price-inelastic demand), market shares converge to a uniform distribution $1/N$, regardless of price dispersion.
4. Higher $\beta$ values amplify market share divergence under identical price differences (in test calibration, share spread expands from 0.089 at $\beta = 0.02$ to 0.632 at $\beta = 0.20$).
5. When demand is well below fleet capacity, served demand matches requested demand exactly without artificial rationing.
6. When aggregate demand substantially exceeds capacity, realized market occupancy approaches 100%.

All verification tests passed successfully on the codebase, confirming the internal consistency of core functional components.

### 3.8 Limitations

Demand parameters ($D_0$, $\beta$, $\sigma_D$) are synthetic and not calibrated against empirical elasticity data due to confidentiality restrictions; findings must therefore be interpreted in **qualitative and comparative** terms across strategies rather than as quantitative projections of realized RevPAR. The choice of proportional rationing (over alternative conventions) is not subjected to sensitivity analysis in the main body of this work, remaining as an avenue for future extension. Furthermore, the model does not account for competitor entry or exit, nor collusive behavior.

# 4 Results and Discussion

## 4.1 Effect of Capacity Constraints and Seasonality

Table 2 summarizes mean RevPAR, its standard deviation across replications, and mean occupancy instability for each strategy-scenario combination, across $N_{\text{seeds}} = 20$ replications of $T = 180$ days each.

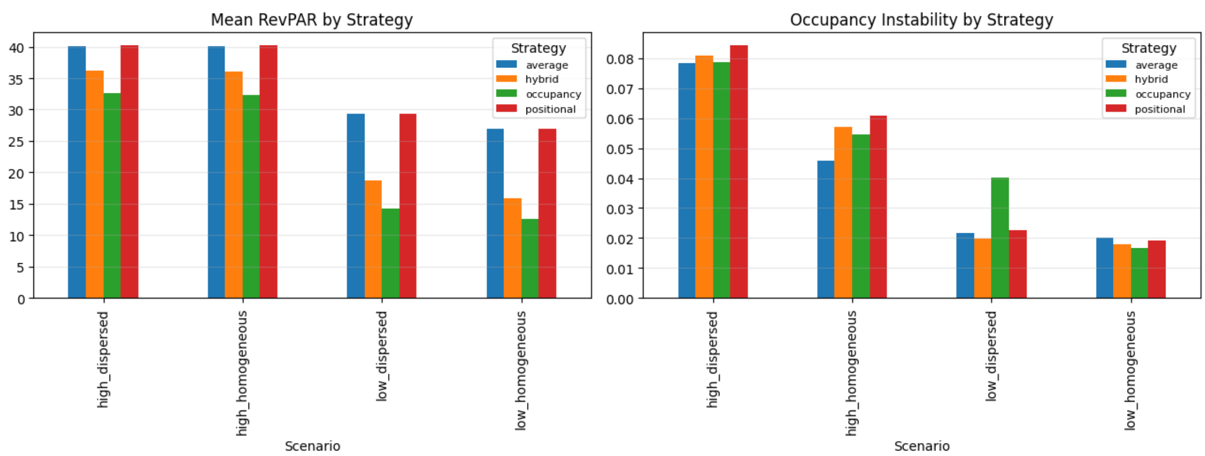


**Figure 2:** *Mean RevPAR (left) and occupancy instability (right) by strategy and scenario. Separate axes by metric.*

| Scenario | Strategy | Mean RevPAR | Std. Dev. | Instability |
|---|---|---|---|---|
| High, dispersed | Positional | 40.263 | 0.945 | 0.084 |
| | Average | 40.144 | 0.666 | 0.079 |
| | Hybrid | 36.239 | 0.791 | 0.081 |
| | Occupancy | 32.551 | 0.550 | 0.079 |
| High, homogeneous | Positional | 40.213 | 0.915 | 0.061 |
| | Average | 40.103 | 0.651 | 0.046 |
| | Hybrid | 36.060 | 0.737 | 0.057 |
| | Occupancy | 32.325 | 0.517 | 0.055 |
| Low, dispersed | Average | 29.364 | 0.481 | 0.022 |
| | Positional | 29.271 | 0.671 | 0.023 |
| | Hybrid | 18.641 | 0.352 | 0.020 |
| | Occupancy | 14.265 | 0.164 | 0.040 |
| Low, homogeneous | Average | 26.987 | 0.506 | 0.020 |
| | Positional | 26.869 | 0.676 | 0.019 |
| | Hybrid | 15.924 | 0.453 | 0.018 |
| | Occupancy | 12.515 | 0.278 | 0.017 |

**Table 2:** *Mean RevPAR (€ per available car-day) and occupancy instability by strategy and scenario, excluding the burn-in period (first 20 days) from metric calculations (20 replications per cell).*

The capacity tension axis generates, as expected, the largest effect on outcomes: mean RevPAR nearly doubles from low to high season across all strategies (e.g., the Average strategy shifts from 26.987–29.364 to 40.103–40.144 € per available car-day), and occupancy instability similarly doubles (from around 0.02 to between 0.05 and 0.08). This result confirms that recalibrating baseline demand $D_0$ (Section 3.2) achieves its intended purpose: generating a market with genuine capacity slack in low season (natural occupancy around 55%) and a tight, yet not fully saturated, market in high season (around

88%).

The fleet size dispersion axis (homogeneous versus dispersed fleets, holding aggregate capacity constant) exhibits a much smaller and more uneven effect: in high season, differences across fleet structures are virtually negligible for all strategies, whereas in low season, the dispersed configuration yields systematically higher RevPAR (e.g., 29.364 versus 26.987 € for the Average strategy). This is likely because greater heterogeneity in fleet sizes allows smaller firms to reach capacity limits sooner and spill residual demand over to larger operators, producing a more efficient revenue redistribution when overall market slack is present.

## 4.2 Q1: Does the Positional Strategy Dominate the Market-Average Rule?

Contrary to initial expectations derived from the empirical observation motivating this study, the *Positional* and *Average* strategies do not differ significantly in any of the four scenarios (Welch's $t$-test, $p > 0.5$ across all cases; see Table 2). The Positional strategy displays a marginal numerical advantage in high season (40.263 vs. 40.144 in the dispersed setting), while the Average strategy shows an equally small lead in low season (29.364 vs. 29.271); however, in both cases, the difference is far smaller than inter-seed variability.

This result is corroborated by a paired-difference test across seeds (Table 3), which leverages the fact that identical market conditions—capacities, initial prices, and demand shocks—are evaluated across all four strategies for each seed. This paired specification provides greater statistical power than independent Welch tests by partialling out common variation from random demand shocks. Even under this more powerful test, no significant differences emerge between Positional and Average across any scenario ($p$ between 0.07 and 0.21), and 95% confidence intervals for mean differences encompass zero in all four cases. The difference comes closest to conventional significance thresholds in the *Low, homogeneous* scenario ($p = 0.070$, mean difference of 0.12 € per available car-day favoring Average), which, while not reaching significance at the 5% level, suggests that a slight advantage for the Average strategy in that specific scenario cannot be entirely ruled out with larger sample sizes.

| Scenario | Mean Difference (Average − Positional) | 95% CI | $p$ (paired) |
|---|---|---|---|
| Low, homogeneous | 0.118 | [−0.011, 0.247] | 0.070 |
| Low, dispersed | 0.093 | [−0.042, 0.227] | 0.165 |
| High, homogeneous | −0.110 | [−0.289, 0.069] | 0.215 |
| High, dispersed | −0.119 | [−0.312, 0.074] | 0.211 |

**Table 3:** *Paired difference test across seeds between Average and Positional on RevPAR (€ per available car-day). No interval excludes zero.*

The paired test across seeds represents the primary inferential specification for this comparison, as it explicitly accounts for the matched simulation design (same seed, identical exogenous conditions); the previously presented Welch test serves as an independent robustness check. Applying the Holm-Bonferroni correction across the four paired tests leaves this conclusion unchanged: with no $p$-values achieving significance prior to adjustment ($p \geq 0.070$ across all four cases), none achieve significance afterward, as step-down corrections never reduce $p$-values.

To evaluate whether the relative value of positional information depends on demand price elasticity, a systematic sweep of parameter $\beta$ (Equation 2) was conducted across seven values ($\beta \in \{0.04, 0.08, 0.12, 0.20, 0.30, 0.35, 0.50\}$), holding all other model specifications—including $N = 10$, $D_0$, and the four capacity and fleet dispersion scenarios—constant, over 20 replications per combination. Figure 3 and Table 4 display the RevPAR difference between Positional and Average, pooled across all four scenarios, as a function of $\beta$.

| $\beta$ | Positional − Average Advantage | 95% CI | $p$ (vs. zero) |
|---|---|---|---|
| 0.04 | +0.048 | [−0.026, 0.121] | 0.210 |
| 0.08 | +0.004 | [−0.073, 0.082] | 0.911 |
| 0.12 | −0.027 | [−0.109, 0.056] | 0.531 |
| 0.20 | −0.059 | [−0.147, 0.028] | 0.189 |
| 0.30 | −0.078 | [−0.169, 0.013] | 0.095 |
| 0.35 | −0.083 | [−0.174, 0.009] | 0.080 |
| 0.50 | −0.090 | [−0.182, 0.003] | 0.061 |

**Table 4:** *RevPAR advantage of the Positional strategy over Average (€ per available car-day), aggregated across the four experimental scenarios, for seven values of $\beta$ (80 observations per row: 4 scenarios × 20 seeds).*

Contrary to a hypothesis formed during a preliminary stage of this study—which compared two sequential model calibrations that varied simultaneously in $\beta$, $N$, and $D_0$, thereby confounding the effect of $\beta$—the controlled parameter sweep **finds no statistically significant difference** between Positional and Average at any evaluated point of $\beta$ ($p \geq 0.061$ across all seven cases). Furthermore, the numerical trend does not align with initial conjectures: as $\beta$ increases, the point estimate—small and statistically insignificant throughout—shifts slightly in favor of Average rather than Positional. We conclude that **there is no empirical evidence within the evaluated range that positional information advantages depend systematically on demand price sensitivity**; differences observed across earlier uncalibrated runs were likely driven, at least in part, by simultaneous shifts in $N$ and $D_0$.

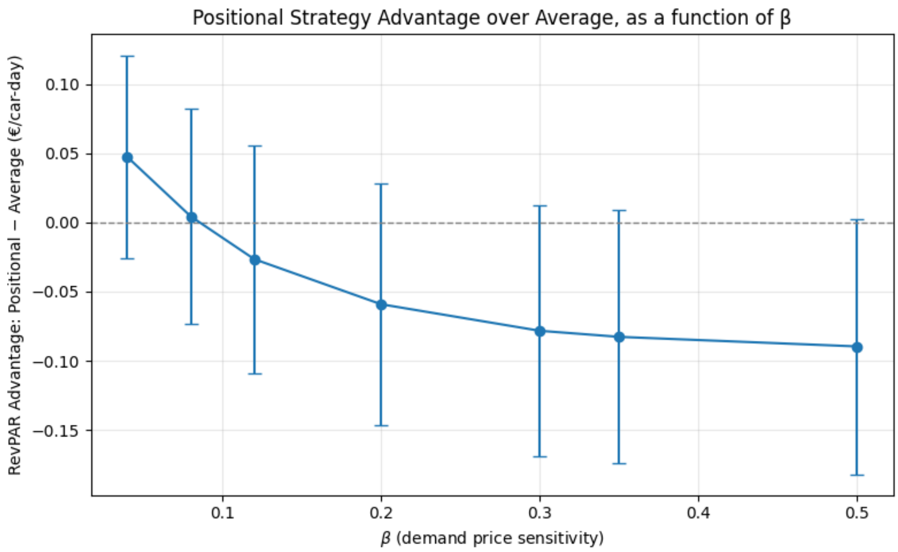


**Figure 3:** *RevPAR advantage of Positional over Average as a function of $\beta$, with 95% confidence intervals. The horizontal line indicates zero advantage.*

## 4.3 Q2: Does the Positional Strategy Reduce Occupancy Instability Relative to Own Occupancy?

Unlike RevPAR results (Section 4.2), which remain stable after excluding the initial burn-in period from metric calculations (see methodological note in Section 3.3), hypothesis tests on occupancy instability **reveal no robust pattern attributable to the choice of pricing rule**. Table 5 summarizes statistically significant comparisons ($p < 0.05$) across the four scenarios.

After applying the Holm-Bonferroni correction for multiple testing (Holm, 1979 [9])—necessary given that 24 total instability tests are conducted, creating false positive risks at nominal 5% levels—most differences that appeared significant under unadjusted Welch tests lose statistical significance. Only two pairwise tests survive correction, both involving the Occupancy strategy in low season (Table 5):

| Scenario | Comparison | $p$ (Welch) | $p$ (Holm) | Direction |
|---|---|---|---|---|
| Low, homogeneous | Average vs. Occupancy | 0.0001 | 0.0021 | Occupancy more stable |
| Low, dispersed | Occupancy vs. others | $< 0.0001$ | $< 0.0001$ | Occupancy less stable |

**Table 5:** *Only occupancy instability comparisons maintaining statistical significance after Holm correction ($\alpha = 0.05$); the remaining 22 of 24 tests lose significance.*

Notably, the two surviving comparisons exhibit an inversion of sign: the Occupancy rule is the most stable under homogeneous fleets, yet significantly the *least* stable under dispersed fleet sizes within the same low-season setting. This pattern—statistically robust post-correction, yet opposing in sign depending on market structure—reinforces the primary conclusion: no single strategy provides general stability advantages, not even Occupancy itself, whose sole robust effect depends critically on market capacity dispersion. All differences involving the Positional rule against Average or Occupancy that seemed significant under unadjusted tests (e.g., in *High, homogeneous*) become indistinguishable from simulation noise once corrected.

This finding rectifies an earlier preliminary observation: an initial assessment, conducted prior to excluding the burn-in period, suggested a systematic stability advantage for the Positional rule in low season. That conclusion does not hold once transient bias is removed: the initial fleet-filling period (where occupancy begins at zero and rises monotonically over early days) artificially inflates temporal occupancy variance unevenly across strategies, creating a spurious difference that vanishes once steady state is isolated. This correction underscores the practical importance of filtering out the startup phase (Section 3.1).

Overall, these findings demonstrate that **within the explored domain, occupancy stability does not reliably differentiate among the four evaluated pricing heuristics**: differences achieving statistical significance are small in absolute magnitude (ranging from 0.001 to 0.02 on a 0-to-1 occupancy scale) and fail to replicate consistently across market scenarios. The operational advantage perceived during industry practice appears, according to this model, to stem primarily from gross revenue generation (Section 4.2) rather than structural inventory stabilization.

## 4.4 Q3: Does the Hybrid Strategy Dominate Both Pure Heuristics?

The most decisive and robust finding across the simulations is that, under the adopted equal weighting ($\gamma = 0.5$), the Hybrid strategy **does not** dominate the two pure heuristics from which it is derived. While Hybrid significantly outperforms Occupancy across all four scenarios (Welch's test, $p < 0.0001$ in all cases), it is surpassed with equal significance by Positional (and, equivalently, by Average) across all conditions. That is, with equal weighting between internal occupancy signals and relative positional pricing, Hybrid consistently achieves an intermediate performance level **between** the worst and best pure strategies, failing to match the performance of the top-performing benchmark.

This result is noteworthy because it challenges the assumption that integrating additional informational signals inherently enhances decision outcomes. In practice, composite heuristics can underperform if one constituent signal exhibits poor predictive efficacy or receives disproportionate weight. This evidence indicates that incorporating internal occupancy signals **depresses the overall performance** of positional tracking under the $\gamma = 0.5$ baseline: the standalone Occupancy heuristic consistently delivers the lowest RevPAR across all configurations (Table 2), and blending it into the Hybrid rule pulls joint returns below what Positional achieves independently. This pattern is consistent with the hypothesis that occupancy feedback operates with an informational lag relative to competitive pricing—reacting to realized utilization from past tariffs rather than current market placement—rendering it structurally less responsive in daily competitive environments.

## 4.5 Synthesis

These simulation outcomes address the primary research question (Section 1) with greater nuance than suggested by early operational observations. Capacity tension and demand seasonality fundamentally dictate aggregate achievable RevPAR levels, but do not alter the relative hierarchy among strategies: the ordering Positional $\approx$ Average $>$ Hybrid $>$ Occupancy remains strictly invariant across all four experimental scenarios. The conjectured superiority of positional rules over aggregate market tracking is not confirmed as a general advantage across seasonality or fleet dispersion regimes, but appears tied to factors outside the original experimental axes, such as price elasticity (Section 4.2). The most conclusive outcome is that pricing solely on internal occupancy yields the weakest revenue performance, and blending it with competitive positioning fails to equal pure positional benchmarking.

Regarding occupancy stability, applying Holm multiple testing corrections leaves only one surviving pattern related to the Occupancy heuristic, which reverses direction based on fleet size dispersion (Section 4.3); no robust, general hierarchy among strategies exists along this dimension.

# 5 Conclusions

## 5.1 Main Findings

This thesis originated from an operational observation during an internship within the *Pricing and Revenue* department of a car rental firm: that monitoring a localized subset of price-adjacent competitors provided a more actionable operational signal for daily rate updates than tracking broader market averages. The central objective was to formalize this premise through an Agent-Based Model, comparing four pricing rules in a controlled synthetic environment across varying capacity constraints and seasonal demand regimes.

The results only partially substantiate initial intuitions, introducing important qualifications. In terms of revenue (RevPAR), no empirical evidence supports the hypothesis that the Positional strategy outperforms market-average tracking across any of the four experimental scenarios (Section 4.2); a systematic parameter sweep further indicates that positional advantages do not emerge systematically across varying levels of demand price elasticity within the tested range. Regarding occupancy stability, once early transient filling biases are removed, no systematic superiority or robust hierarchy across strategies can be identified (Section 4.3): differences remain marginal in absolute scale and reverse direction depending on market fleet concentration.

The most robust finding, however, differs from the initial research proposition: the Hybrid rule fails to dominate either of the top two revenue-generating strategies (Positional and Average), despite incorporating positional competitive signals (Section 4.4). This result—supported by $p < 0.0001$ across all eight relevant RevPAR comparisons—demonstrates that combining heterogeneous information sources does not ensure performance gains over specialized standalone rules.

Finally, capacity tension and seasonality primarily drive the overall level of economic returns achievable by all strategies, roughly doubling RevPAR from low to high season. In contrast, the choice of pricing heuristic substantially impacts relative earnings (with rules ignoring competitor rates lagging significantly behind), but plays a minimal role in dampening aggregate inventory volatility.

## 5.2 Practical Implications for Revenue Management Analysts

From an operational standpoint, this research yields three practical insights for pricing analysts. First (Section 4.2): the lack of significant revenue divergence between Positional and Average rules suggests that tracking a Top-3 cluster of price-adjacent competitors can offer a **more interpretable and operationally tractable** heuristic without penalizing commercial performance—even though sorting algorithms are computationally equivalent to computing arithmetic means, focusing analyst attention on immediate neighbors carries **no observable opportunity cost in RevPAR** under moderate elasticity settings. Second (Section 4.2): these findings provide a **quantitative evaluation** of an industry rule of thumb often applied through intuition, contrasting it with complex multi-signal algorithms that require broader data ingestion without delivering guaranteed incremental revenue. Third (Section 4.4): the empirical evidence demonstrates that arbitrary combinations of occupancy and competitor signals can active degrade returns relative to pure competitive positioning.

## 5.3 Limitations

The limitations of this research stem primarily from deliberate scoping decisions designed to keep the study tractable within the schedule of a Master's thesis, and can be categorized into three areas.

First, structural model assumptions, as outlined in Section 3.8: proportional rationing of spillover demand represents one convention among several theoretical alternatives; vehicle marginal operational costs are omitted, restricting conclusions to top-line RevPAR rather than net profit margins; and rental durations are modeled deterministically.

Second, constraints related to synthetic parameter calibration. The simulation highlighted sensitivity to two baseline parameters—baseline demand $D_0$ and logit price sensitivity $\beta$—which were calibrated using documented criteria (Section 3.2) but could not be validated against proprietary industry datasets due to non-disclosure restrictions. Consequently, all findings should be interpreted in **qualitative and comparative** terms rather than as absolute real-world revenue forecasts.

Third, experimental scope limitations: variations were restricted to two factorial axes (capacity and seasonality), without incorporating elasticity as a formal third factorial dimension or optimizing the hybrid blending parameter $\gamma$. The model also excludes market entry, exit, collusive dynamics, and adaptive strategy switching across simulated horizons. These avenues are detailed further in Section 6.

# 6 Future Lines of Work

The scoping limitations discussed in the preceding section are not primarily model deficiencies, but rather deliberate boundaries established to keep this Master's thesis feasible within the available timeline. This section develops six specific avenues for extension, arranged approximately in ascending order of implementation complexity, proposed as a natural continuation of this work.

## 6.1 Demand Price Elasticity as a Third Experimental Axis

An analysis conducted during the development of this study ($\beta$ parameter sweep between 0.04 and 0.50, Section 4.2) found no evidence that the relative advantage of the Positional strategy depends on demand price elasticity, contrary to a preliminary hypothesis derived from uncalibrated model comparisons. However, this sweep examined a single parametric axis ($\beta$) while holding $N = 10$ and baseline capacity configurations fixed; exploring whether this independence holds across markets with different competitor counts remains a natural extension, particularly given that the discarded preliminary conjecture arose from comparing settings with different values of $N$. Furthermore, the sweep could be extended to $\beta$ values outside the explored domain, or integrated factorially with the capacity and fleet dispersion axes evaluated in the primary design, rather than aggregating findings across the four scenarios as done here.

## 6.2 Sensitivity of the Hybrid Strategy to Parameter $\gamma$

The Hybrid strategy was evaluated under an equal weighting ($\gamma = 0.5$) between own-occupancy and positional signals (Equation 11), without testing alternative allocations. Because this work demonstrates that the occupancy heuristic is, in isolation, the weakest-performing rule across all scenarios (Section 4.4), it is relevant to evaluate whether assigning a smaller weight to that signal ($\gamma < 0.5$) allows the Hybrid heuristic to approximate the returns of pure Positional tracking, or whether any non-zero reliance on internal utilization degrades performance proportionally. A systematic sweep across $\gamma \in \{0, 0.1, 0.2, \ldots, 1\}$, holding all other structural specifications invariant, would trace the full RevPAR response curve as a function of $\gamma$, identifying whether an interior optimal weighting exists or whether the trajectory is monotonic. This extension carries the lowest implementation cost among the six proposals, requiring only the parameterization of $\gamma$ across the established simulation codebase (Section 3.3) without altering the core simulation logic.

## 6.3 Incorporation of Marginal Operational Costs and Net Margin Analysis

The current model evaluates pricing heuristics strictly in terms of top-line revenue (RevPAR), excluding marginal operational costs per rented vehicle (routine maintenance, mileage depreciation, turn-around cleaning, and station handling). A natural progression would assign a synthetic marginal cost $c$ per car-day rented and recalculate performance metrics on net contribution margins (RevPAR minus aggregate marginal fleet costs) rather than gross revenues. This extension is particularly relevant because heuristics that achieve higher utilization—and therefore greater vehicle wear and tear—do not necessarily yield higher unit revenues, meaning that the strategy performance hierarchy (Section 4.5) could shift when evaluated under net operating margins. This step provides the closest alignment with a complete corporate financial evaluation, although it requires introducing an additional assumption (the synthetic marginal cost level) for which reliable public industry benchmarks are scarce.

## 6.4 Heterogeneity in Rental Durations

The model assumes a fixed, deterministic, and identical rental duration ($L = 4$ days, Section 3.2) across all customer cohorts and firms, whereas in practice rental length varies considerably across commercial segments (weekend leisure, weekly travel, long-term corporate rentals). Replacing $L$ with an empirical probability distribution (for example, a mixture of geometric or log-normal distributions calibrated to realistic duration profiles) would introduce secondary variance into fleet return flows (Equation 4), potentially influencing steady-state occupancy stability across heuristics. Implementing this is technically direct—replacing the deterministic return queue 'history_b' in the simulator with a stochastic return-date assignment generated at booking confirmation—though it requires updating the fleet availability accounting structure.

## 6.5 Alternative Demand Rationing Conventions

This work assumes proportional rationing for unserved spillover demand (Section 3.2) without testing against the alternative convention of "efficient" consumer-surplus rationing, despite theoretical results showing that capacity-constrained equilibrium outcomes are sensitive to rationing assumptions (Davidson and Deneckere, 1986 [3]; Section 2.1). Re-running the full experimental design under both rationing regimes would verify whether the strategy hierarchy identified in this thesis (Section 4.5) remains robust to this specification choice or whether it partially reflects the chosen rationing rule. Because both rationing approaches are well documented in the theoretical literature, this extension would also tie the simulation results more closely to the theoretical framework in Section 2.

## 6.6 Calibration with Empirical Data and Endogenous Strategy Adaptation

Over a longer horizon, and provided corporate non-disclosure constraints can be partially relaxed—such as through aggregated or anonymized commercial records—the extension with the greatest applied value would involve calibrating demand parameters ($D_0$, $\beta$, $\sigma_D$) directly from observed transaction data, replacing the synthetic calibration in Section 3.2 with empirical econometric estimates. Along the same lines, given that no single heuristic dominates across all market conditions (Section 4.5), an avenue of academic and operational value would be to allow simulated operators to switch strategies adaptively—such as through a performance-driven selection mechanism based on recent profitability, following heterogeneous agent models with evolutionary strategy switching (Section 2.3). This would expand the model from a static benchmark across uniform market regimes into an inquiry regarding which pricing heuristic emerges endogenously under competitive selection.

# References

[1] Talluri, K.T. and van Ryzin, G.J. (2004). *The Theory and Practice of Revenue Management*. Springer.

[2] Kreps, D.M. and Scheinkman, J.A. (1983). Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes. *The Bell Journal of Economics*, 14(2), 326–337.

[3] Davidson, C. and Deneckere, R. (1986). Long-Run Competition in Capacity, Short-Run Competition in Price, and the Cournot Model. *The RAND Journal of Economics*, 17(3), 404–415.

[4] Simon, H.A. (1955). A Behavioral Model of Rational Choice. *The Quarterly Journal of Economics*, 69(1), 99–118.

[5] Geraghty, M.K. and Johnson, E. (1997). Revenue Management Saves National Car Rental. *Interfaces*, 27(1), 107–127.

[6] Carroll, W.J. and Grimes, R.C. (1995). Evolutionary Change in Product Management: Experiences in the Car Rental Industry. *Interfaces*, 25(5), 84–104.

[7] Tesfatsion, L. (2006). Agent-Based Computational Economics: A Constructive Approach to Economic Theory. In L. Tesfatsion and K.L. Judd (eds.), *Handbook of Computational Economics, Vol. 2: Agent-Based Computational Economics*, 831–880. Elsevier/North-Holland.

[8] LeBaron, B. (2006). Agent-Based Computational Finance. In L. Tesfatsion and K.L. Judd (eds.), *Handbook of Computational Economics, Vol. 2: Agent-Based Computational Economics*, 1187–1233. Elsevier/North-Holland.

[9] Holm, S. (1979). A Simple Sequentially Rejective Multiple Test Procedure. *Scandinavian Journal of Statistics*, 6(2), 65–70.

[10] McFadden, D. (1974). Conditional Logit Analysis of Qualitative Choice Behavior. In P. Zarembka (ed.), *Frontiers in Econometrics*, 105–142. Academic Press.

[11] Anderson, S.P., de Palma, A. and Thisse, J.F. (1992). *Discrete Choice Theory of Product Differentiation*. MIT Press.

[12] Calvano, E., Calzolari, G., Denicolò, V. and Pastorello, S. (2020). Artificial Intelligence, Algorithmic Pricing, and Collusion. *American Economic Review*, 110(10), 3267–3297.

# A Details of Work Development

| Task | Time (hours) |
|---|---|
| Literature review and theoretical framework | 25 |
| Formal model specification | 20 |
| ABM simulator development in Python | 40 |
| Scenario design and execution | 20 |
| Results analysis and statistical testing | 20 |
| Thesis report drafting | 25 |
| Total | 150 |

**Table 6:** *Approximate time dedicated to the thesis*

# B Code Availability and Replication

To ensure complete computational reproducibility, the entire Python codebase of the Agent-Based Model, including the simulation engine, experimental pipelines, calibration scripts, and visualization notebooks, is publicly available in the following GitHub repository:

`https://github.com/Daniseralz/TFM-Dynamic-Pricing-ABM`

The repository includes instructions to replicate all 320 baseline simulation runs, the calibration checks ($\beta$ and $D_0$), and the sensitivity parameter sweeps presented in this work.