Lisbon School of Economics & Management
Universidade de Lisboa

MASTER
DATA ANALYTICS FOR BUSINESS

MASTER´S FINAL WORK
INTERNSHIP REPORT

APPLIED ALGORITHMS AND BUSINESS INTELLIGENCE DASHBOARDS FOR REVENUE DECISION-SUPPORT

BARIS RODI BAYAR

Lisbon School of Economics & Management
Universidade de Lisboa

JANUARY - 2026

Lisbon School of Economics & Management
Universidade de Lisboa

MASTER
DATA ANALYTICS FOR BUSINESS

MASTER´S FINAL WORK
INTERNSHIP REPORT

APPLIED ALGORITHMS AND BUSINESS INTELLIGENCE DASHBOARDS FOR REVENUE DECISION-SUPPORT

BARIS RODI BAYAR

SUPERVISION:
PROF. ARTUR CUNHA

JANUARY - 2026

To my parents Sultan Bayar
and Vahap Bayar who
supported me in every way
possible during my journey
of education and life.

GLOSSARY

1P – First-party (company-owned vehicles/business line)

AI – Artificial Intelligence

Ancillaries / take rate – Extras and their attach rate per booking/night

CI – Check-in (pickup) context/depot/date

CO – Check-out (drop-off) context/depot/date

Conversion – Funnel conversion (search to purchase)

CPT – Commercial, Planning and Trading

CSV – Comma Separated Values

Dashboard – Business intelligence monitoring layer

Data Governance – The adoption, implementation, and monitoring of data ownership, policies, regulations, and procedures across an organisation's data management operations

DE – Date Entered (booking creation date) logic

Decision-support system/tools – Algorithms + dashboards assisting decisions

Depot – Operational location/station

KPI – Key Performance Indicator

MCDM – Multi-Criteria Decision-Making

Minimum nights rule – Minimum length-of-stay rule

Offer rate – Funnel rate from search to offer

One-way rental – Pickup and drop-off at different depots

OW – One-Way

PU – Pickup date (travel start/travel date) logic

SQL – Structured Query Language

Stockout – Situation in which no sellable vehicles remain

Vehicle Class/Vehicle Category – Class is a broader grouping, a class can contain more than one vehicle category

WBR – Weekly Business Review

WoW – Week-on-Week

YoY – Year-on-Year

YTD – Year-to-Date

ABSTRACT, KEYWORDS AND JEL CODES

This report outlines a practical approach to developing decision-support tools for the fleet-based rental businesses, with a specific focus on the revenue management for vehicle fleets. It describes the creation of the algorithmic systems and the dashboards aimed at optimising the revenue by enhancing the visibility, ensuring consistent decision-making and facilitating operational monitoring. The internship project centred on integrating the algorithmic logic with the business intelligence tools to support the decision-making within a global fleet operation, while emphasising the importance of the human oversight and governance.

The project involved three key algorithmic interventions: adjustments to the one-way pricing, strategies for the discount allocation and scarcity tagging, and rules for the search ranking. These tools were designed to optimise the fleet utilisation and pricing strategies, taking into account the fluctuating demand across the various depots and vehicle categories. The development of the dashboards offered the stakeholders a clear view of the operational outcomes, enabling the real-time monitoring and adjustments. The resulting decision-support systems, designed with scalability in mind, have contributed to improving the revenue management practices at Indie Campers.

Challenges arose during the project, including issues with the data governance, time limitations and restrictions due to the unavailability of the sensitive data. Despite these challenges, the tools delivered immediate value by enhancing the decision-making consistency and providing greater operational insight. Future efforts should focus on refining the algorithms for broader application across different locations and further expanding the manual input systems to reduce the possibility of human error.

KEYWORDS: Revenue Management; Fleet Optimisation; Algorithmic Decision-Support; Business Intelligence; Dashboard Development; Dynamic Pricing

JEL CODES: M31, C61, D81, L91, C63

TABLE OF CONTENTS

TABLE OF FIGURES

TABLE OF TABLES

ACKNOWLEDGEMENTS

My sincere gratitude goes to my thesis supervisor, Prof. Artur Cunha, for his guidance, support, and valuable feedback throughout this work. Appreciation is also extended to my professors for their contribution to my academic development during my studies.

Thanks are due to my colleagues and supervisors during my internship for their support and practical insight, which were essential to this thesis.

I am grateful to my classmates for their discussions and shared ideas along the way, and to my close friends for their encouragement and support, especially during challenging moments.

Finally, my deepest thanks go to my parents and my sister for their constant support, patience, and belief in me. I am truly grateful for everything they have done.

APPLIED ALGORITHMS AND BUSINESS INTELLIGENCE DASHBOARDS FOR REVENUE DECISION-SUPPORT

By Baris Rodi Bayar

1. INTRODUCTION

Fleet-based rental businesses operate with capacity that is most of the time fixed in the short term. Since the product is physical, the limitation is not just a theoretical limitation. The inventory comprises a finite fleet, and decisions take shape based on the simple reality of “what gets sold today determines what can be sold tomorrow”.

Because bookings occupy a vehicle over a period of time, the constraint expands from being about how many vehicles exist to when each one becomes blocked and when it becomes available again. This creates continuous tension between monetising demand that is achievable now and retaining sufficient flexibility to serve later demand, which remains unknown but still having the potential to be more valuable.

When operations span into multiple depots, the decision environment transforms drastically because availability and demand across locations rarely stay aligned. Fleet distributions can fall far from demand patterns, sometimes slowly and sometimes abruptly, depending on the realised mix of bookings. The decision made for a depot has the possibility of indirectly shaping the outcome in another, even when those two depots are managed separately on a daily basis.

Revenue decisions in these types of settings go beyond choosing a price level. They include prioritisation choices, the way availability gets used, and the way the booking experience frames value and urgency, since customers respond to what they see on the interface as well as to the underlying offer.

Algorithmic decision-support becomes relevant here as a way of turning repeated judgement requirements into a more consistent logic for decision-making (Sele & Chugunova, 2024). A well-designed tool does not remove responsibility from humans, it reduces arbitrary variation between different stakeholders and offers a stable reference point when they need to align (Nadj et al., 2020).

Dashboards support the same objective by making the outcomes visible and comparable over time or another basis like locations (region/country/depot) or customer segments (international and domestic). Without the monitoring capability, an algorithm becomes difficult to govern and to discuss, since teams end up debating opinions instead of observing outcomes (Sele & Chugunova, 2024; Nadj et al., 2020). Dashboards provide shared visibility and allow the organisations to see whether intended effects are occurring or whether side effects are emerging or not (Park et al., 2022).

Due to confidentiality requirements, this report avoids sensitive figures and highly detailed internal operational information. The discussion will be kept at a level that will preserve business privacy while still documenting the applied methods and deliverables.

The present internship report documents applied work, focusing on practical decision-support tools like algorithms that produce usable outputs or recommendations and dashboards that support interpretation, validation, and controlled intervention. The emphasis is on usability, interpretability, and operational integration, with full automation and theoretical optimality treated as outside the scope for the internship deliverables.

The following chapters first establish the relevant literature and conceptual foundation; then describe the organisational context, methods, data perspective, implemented tools, outputs, and the evaluation logic used to assess outcomes within the boundaries of the internship.

2. LITERATURE REVIEW & CONCEPTUAL FOUNDATIONS

2.1. Revenue Management and Pricing in Asset-Intensive Businesses

In the asset-intensive services, revenue management can be comprehended as managing the limited capacity under uncertain demand whereas still making decisions that are operationally executable (Bandalouski et al., 2018). Moreover, in hotel-focused discussions, revenue management is framed as a practice that links demand with available capacity through pricing and capacity control, it is also treated as an information-system supported discipline instead of a purely theoretical optimisation topic (Bandalouski et al., 2018).

One of the main difficulties in these contexts is that capacity is being consumed over time and cannot be recovered once the service window has started, this supports the historical development of revenue management in the industries with time-perishable inventories (Bandalouski et al., 2018). At the same time, the digital environment has altered market conditions because firms can adjust prices with lower effort and customers can compare offers with lower search cost (Bandalouski et al., 2018).

Car rental adds another network dimension because the inventory of the business moves and bookings are tied to CI (check-in) and CO (check-out) structures across time and location (Li et al., 2023). Vehicles can be transferred amongst stations through customer one-way travel behaviour with the duration of the rental length, creating spatial and inter-temporal demand correlations that complicate capacity assessment (Li et al., 2023). Under these conditions, local decisions become coupled; since the system behaves as a network, not isolated depots (Li et al., 2023).

Alongside these models, it is emphasised that dynamic pricing practices are shaped by information availability and the wider online ecosystem in the hospitality industry (Ampountolas, 2016). Social media and online channels are discussed as changing how consumers compare offers and how they interpret the prices, changing the behavioural response part of the revenue problem (Ampountolas, 2016).

A practice-oriented view of short-term rentals reinforces a similar point from another direction: Revenue management adoption can remain uneven, by strong emphasis on pricing and competition while systematic use of customer behaviour and historical learning may be less mature except in more professionalised settings (Chalupa & Petricek, 2025). This suggests that, in real organisations, the value of revenue tools often comes from improved decision consistency and continuous monitoring, not from introduction of the complex optimisation that stakeholders do not absorb (Chalupa & Petricek, 2025).

2.2. Algorithmic Decision-Support Systems

Prescriptive decision-support is often presented as the part of analytics that goes beyond describing what happened and starts assisting decisions on what to do next, making it organisationally attractive but equally sensitive to the implementation conditions that differ according to the industry (Wissuchek & Zschech, 2025). From an information systems perspective, prescriptive analytics systems can be described as a mature form of business analytics, where strong emphasis should be placed on the relationship between the system and the user, particularly on how the decision authority and responsibility is being delegated (Wissuchek & Zschech, 2025). The design space can be formalised using system archetypes by distinguishing advisory, executive, adaptive and self-governing forms; these describe how much the system acts, versus how much it supports (Wissuchek & Zschech, 2025).

In many applied environments, decisions cannot be reduced to a single criterion without losing important operational meaning, this is why multi-criteria logic appears repeatedly, even when utilised informally (Sahoo & Goswami, 2023). MCDM (multi-criteria decision-making) methods are the tools for complex decisions involving multiple criteria, including families of approaches that incorporate uncertainty handling, hybrid structures, and data-driven developments; that can be used for various scenarios such as ranking and scoring approaches used for prioritising actions when revenue, utilisation, availability and operational feasibility must be taken into account altogether (Sahoo & Goswami, 2023).

Adoption remains a behavioural and organisational question, many decision-makers still doubt whether MCDM tools are beneficial and may prefer intuitive decision-making (Ishizaka & Siraj, 2018). Survey evidence shows that there are possibilities of having a gap between awareness of MCDM methods and actual practice, whereas, experimental evidence indicates that when participants revised their rankings, they followed the tool’s recommendation (Ishizaka & Siraj, 2018). This implies that structured support can influence decisions once users engage with it (Ishizaka & Siraj, 2018).

AI and prescriptive analytics support the applied decision processes in supply chain resilience systems in volatile conditions (Smyth et al., 2024). Related research and practice remain fragmented across technologies and contexts, but even when the domain differs the implication is transferable: Prescriptive support gains value when uncertainty and operational constraints dominate, but it should be framed as a socio-technical system rather than a purely computational achievement (Smyth et al., 2024).

2.3. Behavioural Pricing and Scarcity Signals

Pricing decisions operate in a behavioural environment, especially in online funnels where the customers make their decisions under comparison pressure and attention constraint (Ampountolas, 2016). Scarcity and urgency cues are discussed as mechanisms that can shift perceived value and accelerate choice, even if the underlying product remains unchanged (Zamfir, 2024).

Urgency influences perceived value through behavioural mechanisms and cognitive shortcuts, while scarcity can trigger decision biases beyond rational evaluation (Zamfir, 2024). This framing is important for applied design because it suggests that the scarcity cues function as the interventions that can steer behaviour, creating a requirement for its governance and monitoring (Zamfir, 2024).

Empirical work on time scarcity in live-streaming e-commerce, supports the point that time-based limitation can heighten perceived urgency and impulsive buying, also treats perceived urgency as a mechanism that helps explain the behavioural effect (Hao & Huang, 2023). Product-type effects appear in parts of the mechanism, indicating that urgency cues can be sensitive to context (Hao & Huang, 2023).

Product scarcity effects strengthen this caution by showing that scarcity effects differ by scarcity type and product conditions, a meta-analysis reports (Ladeira et al., 2023). Results indicate stronger purchase likelihood under excessive demand scarcity compared to restricted supply scarcity and under variety scarcity compared to category scarcity; on the other hand, urgency-type scarcity defined through limited quantity and limited time is not supported in the same consistently positive way (Ladeira et al., 2023). Because it warns against interpreting all scarcity signals equally and promotes a more restricted use of scarcity cues that is in line with the organisational context and customer expectations, this heterogeneity is crucial for applied decision-support design (Ladeira et al., 2023).

Behavioural scarcity discussions also take into consideration that such tactics can carry long-term risks related to trust, loyalty and regret; therefore, scarcity interventions should be framed as adjustable levers (Zamfir, 2024). In an applied revenue system, this implies a need for monitoring dashboards and controlled rules, so that the organisation can validate effects and limit unintended results over time (Zamfir, 2024).

2.4. Business Intelligence and Dashboards as Decision-Support

Dashboards are used because decision-making becomes unreliable when information is scattered, inconsistent, or too slow to interpret, especially when the organisation must track many KPIs across time and locations (Park et al., 2022). Visualisation interventions affect decision-making outcomes and perceptions by shaping what decision-makers notice, the way they interpret and the action they take (Park et al., 2022). This process also depends on the mediators such as perceived quality and trustworthiness (Park et al., 2022).

A case study on implementing KPIs (key performance indicators) and designing performance dashboards provides a practical view of the organisational layer: The organisation deployed new KPIs across departments, and the standardised indicators along with the dashboards were used to support managerial analysis and comparison (Nunes et al., 2024). As a result, it was reported that the senior managers were able to make more detailed performance assessments (Nunes et al., 2024). The study results suggest that KPI governance and dashboard design are organisational processes as much as they are technical tasks, supporting the point of consistent definitions and organisational consensus on what is being measured are essential for dashboard efficacy (Nunes et al., 2024).

An experiment shows interactivity can improve decision speed and exploratory power, but it can also produce unintended cognitive effects (Nadj et al., 2020). Although interactive analytical dashboard features in operational decision-support systems may improve task performance by what-if analysis, it also may decrease situation awareness, local tasks can be achieved whilst broader contextual understanding can be missed (Nadj et al., 2020). In order to prevent this out-of-the-loop problem risk for applied decision-support, highly interactive features should be paired with practices that preserve contextual understanding (Nadj et al., 2020).

2.5. Data Governance and Data Quality in Decision-Support Systems

Data that is shaped by the operational processes is essential for the decision-support systems' efficacy, but this is frequently disregarded when the focus is centred on algorithmic design. Incomplete, inconsistent or out-of-date input data enables systems to operate free from any clearly visible errors, while generating the outputs that divert the decision-maker rather than enlighten them (Khong et al., 2023). Stemming from this weakness, the data quality is generally defined by a number of criteria like “validity, timeliness, accuracy, completeness, and consistency”, that capture the various forms of failure (Khong et al., 2023). These factors are still indispensable when talking about data quality in AI environments, because models and dashboards usually assume that the input data is trustworthy by default, which increases the cost of quality errors that go unnoticed (Guillen-Aguinaga et al., 2025).

Data quality is the direct determinant of the decision quality more than just a technical matter (Khong et al., 2023). Research shows that most of the respondents regard the data quality as essential for decision-making, they also relate inferior decisions and increased effort spent for verification of the information directly to poor quality (Khong et al., 2023). In those situations, decision-support can paradoxically become the source of the additional work more than a mechanism that reduces uncertainty (Khong et al., 2023).

Data governance functions as the structure of the organisation that prevents quality of the data from deteriorating with time, particularly in the organisations where multiple systems and teams contribute to the same reporting outputs (Bernardo et al., 2024). A governance approach for most situations can be summarised as establishing the explicit ownership of datasets and indicators, the agreed upon definitions of KPIs and the authority over transforming the logic to prevent different teams from operating with the incompatible figures (Bernardo et al., 2024). In the scope of digital transformation initiatives, the governance is even more important, since as the pipelines evolve, source systems change and the temporary fixes can proliferate if the responsibilities remain unclear and the oversight points are absent (Bernardo et al., 2024).

The reliability of decision-support systems therefore depends on the data processes that stay sufficiently stable to be relied upon and transparent enough to be questioned when irregularities arise (Bernardo et al., 2024). Low data quality decreases the confidence in dashboards and their recommendations, especially in the cases of not being able to distinguish the origin of the problem being the data, the metric definition or the model logic (Khong et al., 2023). Confidence for predictive decision-support is constructed through the perceived reliability and credibility of the outputs, alongside their explicability; so, it stems also from sustained consistency, not from technical sophistication alone (Schwartz et al., 2022).

2.6. Human Oversight in Algorithmic Decision-Support Systems

A practical way for integrating the literature above is treating the decision-support system as a pipeline that transforms the operational data into outputs by the algorithms and exposing the outputs through dashboards so that humans can use them to decide and monitor the outcomes (Wissuchek & Zschech, 2025). This flow well aligns with the perspective of the prescriptive analytics systems, which considers system design as a delegation problem and distinguishes archetypes based on decision authorities share among the user and the system (Wissuchek & Zschech, 2025).

The involvement of human-in-the-loop is generally presented as the benefit of the governance; however, empirical evidence shows that the impact on the performance can be ambiguous and sometimes even counterintuitive (Sele & Chugunova, 2024). Putting a human-in-the-loop increased the uptake of the automated recommendations but they were coupled with a decreased accuracy, hence systematic errors were produced instead of corrections, due to the behavioural patterns of users while they monitored and adjusted the recommendations (Sele & Chugunova, 2024). The relevance of this with the applied revenue tools is high because manual overrides need structure (such as clear triggers, monitoring and accountability) on the contrary of being treated as the general guarantee of safety (Sele & Chugunova, 2024).

Integrating humans with the learning systems brings challenges with it such as bias, data issues, and engagement constraints. Human input can be framed as valuable only when incorporated meaningfully (Kumar et al., 2024). Strengthening an intellectually conservative stance: Dashboards and algorithms are tools for decision-making that are integrated into governance, tracked using predetermined KPIs, and utilised under clear human authority (Kumar et al., 2024).

3. COMPANY & INTERNSHIP DESCRIPTION

3.1. Company Overview

Indie Campers is a fleet-based vehicle (mostly vans) rental company, in which the revenue is generated by the utilisation of a finite set of physical assets that must be available at the right time in the right place. Revenue generation is through the rental and sales of first-party vehicles that are owned by the company, and the rental of third-party vehicles that are owned by others but rented by the company. The foundation of the company dates back to the mid-2010s and during the internship the company had around 500 employees worldwide and the operation networks consist of around 100 depots across 3 continents: Europe, North America and Oceania. The network matters because the inventory is limited, distributed across depots, that are moving through the customer trips, and constrained by positioning and operational requirements.

The company operates by design: even when the overall fleet size seems enough, local availability may become unequal, vehicles are distributed around depots, and bookings take up capacity over travel periods. Because depot circumstances, fleet distribution, and the realised booking mix can all affect the outcome of a commercial decision, the revenue problem is rarely a pure price-setting issue.

Bookings are mostly made via a digital funnel, which means that the customer-facing platform serves as both a primary sales channel and an interface for converting price and availability signals into consumer action. Since decisions must be both operationally practical and economically coherent when presented to customers while still being monitorable and governable internally, this adds another level of complexity to revenue work.

3.2. Internship Context

The internship was completed in the 1P (first-party) Revenue team of the company, in an applied context where the analytical work was anticipated to generate results that the stakeholders who make the decisions can utilise directly, it was not a closed research environment. Operational limits, time constraints, data availability, and the practical necessity that products remain comprehensible and maintainable after initial deployment, all influenced priorities.

A significant portion of the task required cross-functional cooperation. Because the descriptions, thresholds, and criteria for success are required to take into account how revenue decisions are really made and how various stakeholders understand the findings, they could not be adjusted solely from a technical standpoint. Because of this, development was done in an iterative manner: early iterations were created, tested against data behaviour and business expectations, and then improved until outputs were reliable enough to be tracked and reused.

By default, this approach placed the internship in a human-in-the-loop setting. Although tools were created to help make judgments and cut down on noise, business owners still had the last say. Instead of being viewed as exceptions, manual adjustment methods were regarded as valid governance procedures.

3.3. Role, Tasks, and Skills Acquired

During the internship, the role was 1P Revenue Intern in the central revenue structure, responsible for applied decision-support systems that integrate monitoring dashboards and algorithmic logic. In order to increase uniformity in the revenue decisions and to provide visibility on outcomes at the depot, market, and network levels, the tasks were divided across multiple interconnected streams.

Algorithmic decision-support was the focus of one task cluster, which included work on one-way decision settings and ranking-type logic that ranks actions according to various operational and commercial signals. Transforming current spreadsheet-style reasoning into repeatable scripted procedures while maintaining the computation's transparency to allow for team inspection, challenge, and maintenance was a recurring demand. Another stream concentrated on consumer-facing marketing strategies like scarcity messaging and discounts. Clear boundaries, careful rule design, and the capacity to track impacts without assuming steady behavioural outcomes were all necessary for the study.

Dashboards and routine monitoring, mostly via Metabase, were the focus of a second task cluster. This included creating focused views for the concerns that frequently came up in revenue talks (such as allocation-related checks, price and length-of-stay viewpoints, and promotion monitoring) and maintaining and enhancing stakeholder-facing dashboards used in daily or weekly review contexts. Additionally, a portion of the internship comprised tiny corrective tasks and routine quality checks, which are crucial for maintaining the dependability of decision-support tools in practice but are not often apparent in a final thesis narrative.

From a skills standpoint, the internship improved the capacity to distinguish between what is desired to measure and what is actually quantifiable, as well as to transform the unclear revenue queries into practical logic under actual restrictions. Practically speaking, it also improved the stakeholder communication by matching the KPI definitions, clarifying why the outputs act as observed, and pointing out the instances where limitations of the data need cautious interpretation. A recurring observation throughout the process was that, decision-support is typically more about keeping the control than it is about maximising automation. This entails creating technologies that improve visibility and consistency while maintaining clear human authority for the high-impact revenue activities.

4. METHODOLOGY

4.1. Data Sources

The internal operational data produced by the booking flow and the fleet availability tracking served as the foundation for the analytical work detailed in this paper. Because booking-level records offer the most direct representation of demand that translated into bookings (including the CI and the CO structure, the travel window, and the depot context in which the decision materialised) they served as the foundation for the research.

The data was evaluated in conjunction with operational availability signals to make the booking information useful for revenue decision-support. This required using fleet-state aspects, which differentiated between cars that are reserved, free, and unavailable because of operational limitations (such as repairs and maintenance). This is significant since outputs frequently depended on whether scarcity was actual or only apparent. Considering availability-based filters can produce deceptive gaps if they are not examined, the same reasoning also necessitated tests for extraction artefacts and completeness problems.

Instead of being viewed as a single figure, the revenue was viewed as a composite. In order to separate the base rental product from other financial factors that can influence the final booking value, such as the applicable promotions and the extras that could be sold beside the vehicle rental, the methodology depended on breaking down the financial outcome into interpretable components. This breakdown was utilised for decision relevance in a number of analyses, where the appropriate inquiry was not the total by itself, but rather which portion of the total moved and under what circumstances did it move.

Since operational limitations and revenue decisions do not change along a single calendar axis, time was also managed across several parallel dimensions. The distinction was regarded as crucial, and both the booking date and the travel dates were taken into account. While travel-date aspects were necessary to assess capacity pressure and operational viability, booking-date views were used to comprehend demand capture and funnel behaviour. Since week buckets reduce noise while maintaining pacing patterns that revenue stakeholders generally watch, a number of dashboards and checks were constructed around weekly aggregations of booking, travel, CI and CO dates.

Lastly, because the decision environment in a networked fleet business is rarely consistent, the data was divided by the product structure and the location. The analysis was typically based on vehicle classifications or product groupings that are relevant for operational planning and customer preferences, and region, country, and depot dimensions were regarded as the primary location filters.

4.2. Tools & Technologies

The methodology combined the dashboarding tools, spreadsheet control layers, and script-based implementation; because, the deliverables required to be applicable for the stakeholders and they needed to be repeatable at a rhythm that fitted the operational cycles.

ChatGPT and Google Gemini AI chatbots were used to reduce the friction of code preparation for the algorithms, Perplexity and ChatGPT AI chatbots were used for the research process of the literature review and the language refinement of some sentences and paragraphs of this report. All of the chatbot outputs were extensively analysed and amended if needed.

Metabase was used as the primary business intelligence environment. It served as the dashboard layer where stakeholders could track KPIs, pacing signals, and algorithm outcomes through filters and consistent views. Additionally, it served as a practical query layer during iterative development. Many times, the core logic was formalised into reusable query outputs and then reused across dashboards to reduce definition drift.

Spreadsheets (Excel and Google Sheets) were used as a control and governance layer. They were relevant when manual inputs were required, when mapping tables had to be maintained, and, when stakeholders needed a transparent format for inspection and adjustments. In practice, spreadsheets also functioned as a bridge between scripted outputs and their operational adoption, because they allow review and structured overrides without requiring stakeholders to run code.

SQL and Python were used for data extraction, transformation, and algorithm implementation. SQL was applied to retrieve and structure data from internal databases in the forms that could feed either dashboards or scripts. Python was used to implement decision-support logic and generate structured outputs. These outputs were organised with reproducibility in mind, since the practical value of an algorithm depends on whether results can be contrasted between the runs and if they can be updated reliably.

4.3. Working Process

The approach involved a practical and cyclical method. In this setting, requirements often change when initial outputs become visible to the viewers, making this applied and iterative cycle quite appropriate. Most delivered items started as an initial version, perhaps a dashboard or a basic script. These would then undergo several feedback sessions, allowing the business interpretation adjustments and the technical corrections to develop simultaneously. This approach was necessary because revenue inquiries often begin without strict limits. But, once real data reveals actual behaviour, metric definitions accompanied by rule boundaries tend to surface.

The validation process involved the systematic reviews and input from various groups. From the technical standpoint, the outputs were reviewed to ensure that they were complete and covered the intended segments. Any unusual breaks were also looked for in the data, that could suggest problems with the extraction process or the discrepancies in definitions. From the business viewpoint, stakeholders reviewed the outcomes to ensure that if the logic was appropriate in operational contexts, if the recommendations were still clear, and if the unusual situations were handled with care. When possible, outcomes were checked against established baselines. This was primarily done to illuminate the relationship between benefit and cost, while avoiding presenting modifications as advancements in the absence of sufficient evidence.

Manual inputs were retained to ensure a clear aspect of the governance framework. In several areas of the work, it was found essential to maintain the ways to make specific adjustments. This was for situations where the standard procedures were not enough, or when local circumstances meant it was needed to allow for exceptions. This involved keeping the input sheets current for overrides, updating mapping lists as operational names changed and planning updates at a pace that stakeholders could reasonably manage. The methodology employed for this work can be summarised as a controlled applied workflow. It began with the establishment of the decision context, which was then converted into quantifiable signals. From there, repeatable logic was implemented. Validation occurred through data verification and assessment by the stakeholders, and any manual interventions were documented to ensure transparency and accountability.

5. APPLIED ALGORITHMS FOR REVENUE DECISION-SUPPORT

This chapter presents three applied algorithms developed to support the revenue decisions through repeatable logic and a dedicated monitoring layer. The algorithms were run in Python notebooks, while SQL was executed inside the notebooks to extract the booking, search and availability data from the database. For the controlled input layer of manual intervention, online spreadsheets were used. The outputs were produced as structured tables in CSV format, that were then imported to the company platform, results to be tracked by Metabase dashboards. Since the environment is networked and decisions taken in one location were able to influence the subsequent availability in another location, human authority remained. The tools were built to stay adjustable and easy to inspect during routine operations. For daily use, dashboards were assembled from stable base-table queries and applied with consistent filters (region, depot, vehicle category and time period) to keep comparisons coherent over time.

5.1. One-Way Pricing

5.1.1. Algorithm

One-way rentals were regarded as a means to stimulate the demand and decrease the depot imbalances, depending upon maintaining a fee structure that aligns closely with the operational pressures. When a van is picked up at a CI depot and dropped off at a different CO depot, the customer is essentially carrying out a relocation. The network may either be supported or disrupted by that relocation according to the chosen direction. Fixed charges introduced repeated discordance where certain paths deterred demand when movement was beneficial, and conversely, other paths compelled relocations that subsequently necessitated costly corrections.

The algorithm was structured as a layered pipeline, which allowed its outputs to remain interpretable and governable (simplified pipeline and governance schematic can be seen in Fig. 1). Initially, an average one-way charge was established by employing the approximate route distances. Distance was categorised into a small number of distance bands, with a final band for the unknown distances, so that the structure remained clear and communicable. The average total nights measure for each CI-CO pair was calculated using the historical one-way rentals, with a threshold for the value to be used or else allocating the zone average total nights. This helped in evaluating the “possible” CI date utilisation at the CI depot, because the one-way rental outputs were uploaded for checkout dates. Additionally, the boundary tables that were kept as modifiable inputs, were used to enforce the minimum and maximum one-way fees, giving stakeholders control without requiring code modifications. The data sources, variables and data preparation steps are detailed in Appendix A1.

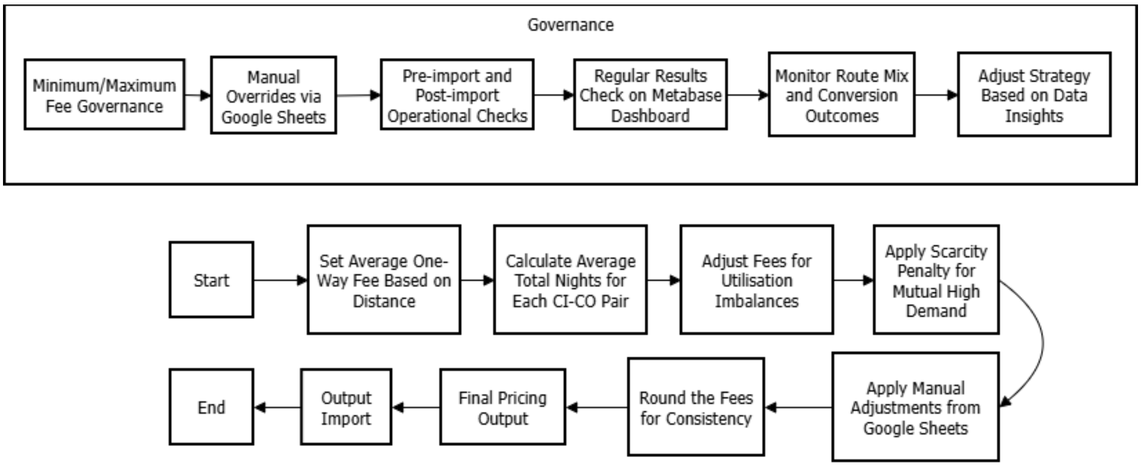

Figure 1 – Simplified One-Way Fee Pipeline & Governance Schematic

An adjustment was then made to account for the pressures arising from utilisation imbalances. The calculation of the utilisation considered a broader travel period and several relevant dates were taken, since a single day's figures would not accurately reflect capacity management, particularly with the one-way bookings because the relocation of the vehicle should be useful in the capacity or revenue aspects. The pressure on the CO side was assessed over a period shifted forward, corresponding to the time when the vehicle will be available and might be needed. In contrast, the pressure on the CI side was assessed around the anticipated CI period (during which the vehicle might be needed already on the CI side on the anticipated CI date) and also around the anticipated CO period (during which vehicle might be needed already on the CI side on the CO date). The average nights approximation established the placement of these windows. These pressure data served to adjust fees, increasing them when CI side pressure (at CI and CO dates) was higher and decreasing them when CO side pressure (at CO date) was higher. This approach was adopted because the same relocation had distinct network meanings, depending on the utilisation of specific dates and its specific origin and destination points. In practice, high CI side pressure was seen as the stricter constraint in this context, because removing a vehicle from a tight origin depot raises the immediate risk of stockout, whereas not having a vehicle in the destination location has a lower risk of lost revenue.

When gaps were found in the information or only had parts of the data at certain times, these situations were approached as simply how the data naturally occurred and fallbacks were included to ensure that every CI-CO combination still produced an output. When utilisation was missing on one side, the adjustment was dampened and the final prices were closer to the average prices, rather than allowing a one-sided/missing signal to dominate. During the critical stages, intermediate tables were printed for review to ensure fallback procedures worked as expected prior to proceeding. This layer lowered the chance of a run appearing to succeed and quietly allowing for poor results.

A second adjustment layer addressed absolute scarcity. Even when a direction looked favourable, the network became fragile when both CI and CO depots approached stockout for a vehicle category. In that setting, encouraging the one-way flow increases disruption risk, and therefore, a scarcity penalty was applied to avoid over-incentivising the routes when the free vehicles on both sides were limited.

Lastly, rounding was used. Rounding followed the regulated procedures to maintain uniformity in customer-facing costs. A dense format was created for the output. Vehicles may emerge at depots in between runs, so, all regional vehicles were included in all depots and any missing combinations were filled with the average values; as a result, risk was decreased and missing rows were not mistaken for “0 one-way fee”.

5.1.2. Manual Adjustments

The explicit governance layer of the manual involvement for the one-way fee overrides were made possible through the use of an online Google Sheets file, enabling the CPT (commercial, planning and trading) associates to enter modifications without changing the code. The overrides were designed to target particular CI-CO depot pairs, specified date ranges and a vehicle category or an aggregation of these. Additionally, the sheet supported a conditional override logic where it enabled the users to input certain algorithm implementation periods and certain free vehicle number thresholds for the rule to be considered.

After each algorithm run, the sheet presented the average of the replaced values by each rule, which helped improve the decisions. The average of outputs was placed alongside the existing manual figures. Percent differences between the override and the average were also presented in summary columns. This alteration decreased the chances of overrides being implemented without understanding their situation, providing a practical method to support the human-in-the-loop strategy.

A second spreadsheet layer covered the minimum and maximum fee governance. Where the CPT associates could input the manual minimum and maximum one-way fees per depot combination.

Operational checks were applied before and after the import of the outputs. Pre-import checks confirmed there were no missing values, no unexpected discontinuities between dates and no other obvious anomalies. Post-import checks reviewed if the website was utilising the correct one-way fee values that were imported.

A discrepancy was observed shortly after the initial deployment, when imported values were compared with the website one-way fees. Investigation with the IT team showed that the platform was not reading the same date keys assumed in the export structure. Once the platform’s date logic was clarified and corrected, the export and import alignment were done. In practical terms, fee selection was confirmed to be driven by the CO date. The export coverage was then adjusted so the correct mapping was always present. Without this fix, the notebook outputs could have remained correct while the customer-facing fees were still wrong.

5.1.3. Results

Results were monitored through a Metabase dashboard tracking the one-way share, pricing behaviour, route mix and conversion outcomes. Views were provided by the booking dates and travel dates, since one-way effects can occur differently across these axes. Filters by the region, CI depot, CO depot and vehicle categories were included to enable spotting local irregularities by not letting them hide inside the network averages. The dashboard centred on plots used for routine checks such as the conversion comparison (see Fig. 2), fee distribution comparison, route desirability comparison and the pace of rentals; comparison was divided under the variations such as route type, fee type, travel week and booking confirmation week. Timely comparison was performed through the comparison of the current post-implementation period with the data of last year same period (to reduce seasonality effect) and certain amount of time before the implementation (to have a general view of the last previous days also).

Figure 2 – Conversion Bar Charts from the One-Way Algorithm Dashboard (censored)

Conversion for the one-way bookings was observed to increase by roughly 25% after deployment. This uplift was primarily attributed to the fees becoming more variable and more aligned with what the customers see as reasonable, so the one-way options no longer appeared as a uniformly expensive special case, including the routes where the one-way movement supports fleet positioning. Also, a pattern consistent with international segments was observed, one-way travel tends to be more appealing for international customers, where returning to the same point is less attractive when the trip is not a journey back to one’s home.

Pricing views separated the one-way fee effects from the base rental price movements. A decrease in the average one-way price per night was observed, consistent with encouraging the one-way choice where directional signals supported it. Net margin was stable or improved, depending on the filters applied. This was interpreted as the combined effect of the high-fee bookings persisting on the costly directions, recovering the lower fees on the beneficial routes, alongside the reduction in the need for the paid relocations by the professional drivers during the stockout periods, since some of the repositioning was achieved through the customer one-way travels.

Directional desirability outcomes were tracked through a further diagnostic view and showed a clearer separation between the route types after the implementation of the algorithm, with the desired directions receiving more attractive pricing signals and the non-desired routes being less, resulting in a higher proportion of the one-way travels being booked in the desired direction. The algorithm outputs were reviewed before import, the booked one-way rentals and their one-way fee details were checked daily, and the dashboard was revisited routinely to keep the tool aligned with the operational conditions.

5.2. Discounts & Vans-Left Tags

5.2.1. Algorithm

The discounts and the vans-left tags were implemented as a combined customer-facing intervention. The discounts were treated as a lever to shape the demand into lower-utilised vehicle categories future periods, while the vans-left tags were treated as a scarcity signal designed to reduce the hesitation when the scarcity is meaningful. Since these interventions act through perception, they were kept adjustable, supported by the monitoring views and the manual control points.

A core design choice was the weekly decision rhythm. The travel dates were grouped into week buckets so that the decisions stayed stable within a week and did not react to the daily noise. The bucket boundaries were chosen by using the analysis of the lookahead weekly utilisation behaviour, the future weekly utilisation curves were reviewed to identify where patterns tended to shift, enabling the near-term, mid-term and far-term behaviour to be captured using different buckets. As an illustration, sharper utilisation changes were typically observed in the near-term horizon, while the subsequent weeks moved more gradually, and the bucket breaks were set where those shifts appeared repeatedly considering all locations. As a result, the weekly aggregation reduced noise and supported consistent decisions for each bucket. The data sources, variables and data preparation steps are detailed in Appendix A2.

Depot capacity structure was then incorporated. Depots do not carry the same variety of the vehicle categories, and the number of available categories influences how promotions should be distributed, since there should be a limited number of discounts in every location unless there is a special promotion. The depots were grouped into a number of category-availability buckets based on how many categories were typically available. In the low-variety depots, the discount spread was intentionally tighter, since the customers have fewer alternatives and broad discounting can become confusing or misleading.

Utilisation quartiles were calculated per depot and week bucket -every time the algorithm was run- to avoid overly precise thresholds and to keep the system robust. Each depot-week combination was positioned within its own distribution, and the discount allocation was mapped accordingly. Lower utilisation buckets were prioritised for the discount assignment to increase the fill in the weeks where the inventory would otherwise remain idle. The discount distribution across categories in those depot-utilisation quartiles were controlled so the incentives did not become too broad or very limited and lose meaning (decision table could be seen in Table I). In practice, the number of the discounted categories and the discount intensity varied by depot buckets, high-variety depots were able to carry a wider spread while low-variety depots remained more conservative.

TABLE I – A PART OF THE DECISION TABLE OF THE VANS-LEFT TAGS & DISCOUNTS ALGORITHM (CENSORED)

			Vans Left			Discounts		
Weeks	No. of Vehicle Categories	Quartiles	Y	Y	Y	Y	Y	Y
Y-Z	Y or less	[Min-Q1]	x	x	x	x	x	x
		[Q1-Q2]	x	x	x	x	x	x
		[Q2-Q3]	x	x	x	x	x	x
		[Q3-Max]	x	x	x	x	x	x
	Between Y-Z	[Min-Q1]	x	x	x	x	x	x
		[Q1-Q2]	x	x	x	x	x	x
		[Q2-Q3]	x	x	x	x	x	x
		[Q3-Max]	x	x	x	x	x	x
	Z or higher	[Min-Q1]	x	x	x	x	x	x
		[Q1-Q2]	x	x	x	x	x	x
		[Q2-Q3]	x	x	x	x	x	x
		[Q3-Max]	x	x	x	x	x	x

Vans-left tags were driven by a related, deliberately different signal. The base scarcity indicator was the average number of free vehicles for a depot-category-week bucket. A tie-break rule was applied where similar average free counts existed across multiple categories, the category with higher utilisation was tagged. This prevented the identical free counts from masking the different risk levels and made the tagging more truthful in operational terms. The tag assignment was capped per depot-week, and the rules were stricter when the average number of free vehicles were not low enough to justify scarcity messaging, keeping the tags credible instead of turning them into a default label.

Output completeness was protected through a dense output structure. The defaults were generated for all regional vehicle categories across all depots, since the vehicle categories can appear in the depots between runs due to fleet movement and snapshot timing. The dense outputs reduced the import failures and avoided silent “no value” behaviour for the missing combinations.

Because discounting changed the price directly, a platform-side safeguard was also used. During the import, the minimum price constraints for each vehicle category were checked by the platform. If a discount would bring a price below the minimum, a warning was raised for the importing user and a lower discount level was applied for that depot-vehicle category-week combination. This created a practical governance checkpoint by the surfacing boundary issues at the import time, not relying solely on the subsequent customer-facing review.

5.2.2. Manual Adjustments

Manual adjustments were enabled through an online Google Sheets file allowing the CPT associates to override the discount level and the tag assignment at the depot-vehicle category-week bucket level. This was treated as necessary for the governance, since the local context can shift within a week due to the maintenance blocks, local events or operational changes that were not captured cleanly in the aggregate signals. Before the imports, the override sheet was usually reviewed alongside the latest output data so that the extreme replacements were visible early enough and unintentional carry-over from earlier weeks was avoided.

The sheet displayed the algorithm output next to the manual overrides, after the algorithm was run, reducing the inconsistent edits. Obligatory time-bounds for overrides helped prevent the promotions from becoming permanent by accident, since the promotional settings often persist once inserted unless expiry is enforced.

Cross-functional alignment was used when appropriate. The discount spread and their intensity were reviewed with the stakeholders outside the revenue function, including the growth-related roles, since these interventions shaped the funnel perception and should not be owned by one team alone. Trade-offs were commonly discussed between the near-horizon conversion gains and the longer-horizon price integrity. The manual layer remained available to adjust when those trade-offs shifted by market or season. Dashboards were reviewed routinely, intermediate notebook extracts were checked when any irregularities appeared, and the spot checks were performed after the import on the selected depots and categories to confirm the week buckets and the values displayed as expected on the platform.

5.2.3. Results

Results were monitored through a Metabase dashboard focused on the funnel behaviour. The conversion rate and offer rate (see Fig. 3) were treated as the leading indicators, since the discounts and the tags influence customer behaviour before longer-term revenue effects are visible. The searches and the booking volume were tracked to detect the demand deterioration within the specific week buckets. The price views were used to confirm the price architecture remained intact. The dashboard followed the same week buckets used by the algorithm, and before-after comparisons were anchored around the deployment period to support consistent interpretation. Additionally, using the base-tables created as the base for the dashboards, the before-after comparisons were easily anchored to the newer import dates.

Figure 3 – Sample Dashboard Plots of Conversion and Offer Rates Before and After the Algorithm Implementation (censored)

Conversion was reviewed overall and by each week-ahead bucket. The monitored window did not show deterioration after the deployment and mild improvements appeared in parts of the horizon. This was interpreted as the combined effect of the discounts improving attractiveness in the lower-utilised future weeks, alongside the tags accelerating the decisions in the pressured categories where -from the customer's perception- any delay would likely lead to stockout. The offer rate was tracked in parallel and remained generally stable or slightly improved due to the same reasons.

Price checks indicated that the discounting remained within the commercial boundaries. The average daily price moved moderately while the volume indicators stayed healthy, showing controlled demand shaping. The monitoring remained cautious because the behavioural interventions can carry trust risks. Tags were reviewed through the spot checks, with the removal by manual adjustment being used where the local context indicated a tag would be misleading. Where the platform minimum-price warning was triggered, it was treated as a follow-up signal and the affected depot-category-week cases were checked again in the dashboard to confirm the intended floor behaviour.

5.3. Search Ranking

5.3.1. Algorithm

Search ranking was accepted as a decision-support problem because offer ordering effectively acts as allocation. Even when the prices remain unchanged, ranking determines which vehicle classes/categories receive attention and which ones are relatively hidden, shaping demand capture and utilisation. The ranking logic was positioned at the intersection of the revenue management and the funnel design.

Vehicle classes were prioritised first and later the vehicle categories within each class; by using multiple performance indicators while keeping the output (see Table II) simple enough to review, adjust and import. The booking history was used as a revealed-preference proxy, while the daily values were avoided due to the noise. The recent booking counts were aggregated over rolling windows, with the most recent period carrying more influence, in order to preserve the responsiveness without causing unstable switching behaviour. The data sources, variables and data preparation steps are detailed in Appendix A3.

TABLE II – EXAMPLE PIECE OF THE OUTPUT FILE (CENSORED)

city	van_category	11-16	11-17	11-18
brisbane	oc-active-bunk-4-manual-base	1	1	1
brisbane	oc-active-compact-2-auto-select	x	y	z
brisbane	oc-active-poptop-4-auto-select	z	x	y
brisbane	oc-active-poptop-4-manual-base	1	1	1

Supply-side signals were incorporated using the forward-looking utilisation and availability. Promoting a popular vehicle class/category when it is close to stockout can reduce the offer rate and harm the conversion by surfacing the options that disappear quickly. As a result, it was decided to prioritise lower future utilised vehicles, also to promote the less utilised vehicle classes/categories in addition to preventing stockout. A blended signal coming from higher past booking numbers and lower future utilisation, supported the promotion of the likely-to-convert vehicle classes/categories while maintaining the offer stability. In practical terms, the change functioned as a correction and stability improvement. The primary goal was to remove the repeated edge-case incorrect ordering while keeping a familiar structure for the stakeholders.

The intermediate outputs were reviewed at the key steps. The class-level ranks and the category-level ranks were printed and checked against the expectations before the import. During the internship, it was explained that the North American region vehicles were supposed to have a specific ordering; so, on the day of the algorithm deployment, this required ordering was implemented easily by the usage of the manual adjustments.

5.3.2. Manual Adjustments

Manual overrides were enabled through an online sheet allowing the CPT associates to set a specific sorting order when required. The date ranges, target regions/countries/depots/vehicle classes/vehicle categories were included as possible specifications. This allowed the ranking actions to be applied for a campaign window or operational disruption period and then expire naturally.

Its purpose was to provide a controlled route to enforce a business decision when needed, such as increasing the exposure of a category for balancing reasons or reducing the exposure due to operational limitations. The outputs were reviewed before the import, the spot checks were completed after the import to confirm if the platform ordering matched the intended results or not, and the periodic checks were run on the performance indicators to catch the possible unalignments or any unwanted results.

5.3.3. Results

Ranking effects were monitored through the funnel stability indicators per vehicle classes and vehicle categories across the locations and time periods. The offer rate was treated as a key leading metric, since surfacing specific categories firstly tends to change the number of offers clicked per search. After deployment, the offer rate remained stable in the depots where the number of vehicle categories were already low, so the elasticity of their rankings did not cause a remarkable change; meanwhile, slight improvement was observed in the depots that had enough vehicle categories, that their ranking was significant. The conversion showed a small positive increment in some of the depots, consistent with the customers seeing the offers that remained available long enough to be booked, reducing the funnel friction. The monitoring view used a time-series comparison around the change period.

Monitoring checks remained part of the routine operations. As mentioned earlier, the ranking-related performance indicators were reviewed periodically to detect any irregular drops, and the spot checks were used to confirm the manual overrides and the notebook outputs were reflected in the visible ordering. Since the ranking changes can generate side effects, these checks were treated as the part of the operational deployment.

6. BUSINESS INTELLIGENCE DASHBOARDS

A monitoring layer was maintained alongside the algorithm outputs, since the decision-support tools were treated as systems that require continuous oversight. Also, most business-related questions require an answer depending on the interpretation of the past data, where dashboards were useful. Metabase served as the shared environment, largely because it supported access control, fast iteration, and consistent filtering by region, country, depot and time period. A base-table approach was adopted, the core datasets were curated first, and the dashboards were then built on top of these tables, so that a definition update could be applied once and then reflected across the plots. The verification relied on the targeted spot checks against the source-of-truth reference numbers and on the stakeholder confirmation, with particular attention given to the most business-critical views.

6.1. Weekly Business Review (WBR)

The WBR dashboard was created to replace an Excel-based weekly review file that was heavy, slow and difficult to govern across the organisation. The same WBR logic was transferred into Metabase so that the review could run with a reasonable refresh time and so that access would not depend on opening a very large file.

KPI alignment was treated as the first constraint. A separate week definition logic was applied, since the standard calendar weeks did not match how the WBR comparisons were discussed in practice. DE (date entered) blocks were aligned to Sunday to Saturday weeks when the WBR dashboard was opened on Monday or any later day that week. PU (pickup) blocks followed the Thursday to Wednesday Indie week logic, since the travel pacing discussions already used that rhythm. This kind of weekly aggregation was therefore treated as a governance decision, because the small boundary changes can shift the interpretation.

The layout followed a drill-down ladder. The regional tables were placed in the beginning for scanning, then country tables and the depot tables for deeper checks. When the Metabase pivot-table limits prevented a perfect replica of the previous layout, the row-header pivot-tables were used, since the stability and the usability were more important than the layout.

DE blocks covered the revenue and the growth with the standard comparison logic [WoW (Week-on-Week), YoY (Year-on-Year) and YTD (Year-to-Date) variants]. The international and domestic splits were retained, since the mix changes repeatedly shaped the conclusions. The deals were excluded from the DE Growth views to avoid the misleading jumps that were not driven by the routine trading performance. The PU blocks paired the revenue and the nights travelled, with the cancellation rate and the fleet utilisation context. The ancillaries were handled through the take-rate views aligned to the nights-based logic, so that the capacity consumption and the commercial outcomes remained comparable. The result was a weekly dashboard (see Fig. 4) that could be used by a broad group of stakeholders, with an established meeting rhythm built around it.

DE Revenue - Region

Region Name	Rev #	Rev WoW	Rev YoY	Rev YTD #	Rev YTD YoY	Nights (Exp) #	Nights (Exp) WoW	Nights (Exp) YoY	Nights (Ex
APAC									
Central Europe									
North America									
UKI & Nordics									
West Europe									
Grand totals									

DE Growth COS & Marketing - Country

Depot Ic Region Name	Depot Country Code	COS #	COS Ref. & Com	COS vs Target (	COS WoW	COS YoY	COS YTD #	COS YTD YoY	(
APAC	AU								
	NZ								
Totals for APAC									
Central Europe	AT								
	CH								
	DE								
	GR								
	HR								
	IT								
	PL								
Totals for Central Europe									

Figure 4 – Sample WBR Tables (censored)

6.2. Daily Pricing per Length

This dashboard was created for the evaluation of an existing daily pricing per length algorithm implementation, where longer searched trips were positioned with a lower price per night. For the comparison, the post-implementation period was matched against the same period last year, and the length bucket sizes were kept consistent with the algorithm's own buckets.

Filters were kept practical for daily use (date entered, region, country and depot). The main focus was the conversion behaviour and its comparison, the conversion by length buckets and by time-series views allowed the evaluation of the efficiency of the algorithm in increasing the conversion per length bucket and time period. The searches and the bookings by the length buckets were shown alongside the conversion, since the conversion can shift due to other changes and a recurring stakeholder question was whether the movement reflected the demand dynamics or the pricing effects. The prices were also used, to compare the total base rental revenue and the average daily base rental price per length bucket, before and after the algorithm implementation. The base rental views were treated as the cleaner read for this topic, since the promotions and the extras can distort the totals.

A broadly positive pattern was observed. The conversion improved more clearly in the higher length buckets, while the average daily prices did not show a disproportionate decline, supporting the view that the demand was being shaped in a controlled manner (see Fig. 5). The alignment with the stakeholder responsible for the underlying pricing logic was used to confirm the timeframe assumptions and the edge cases.

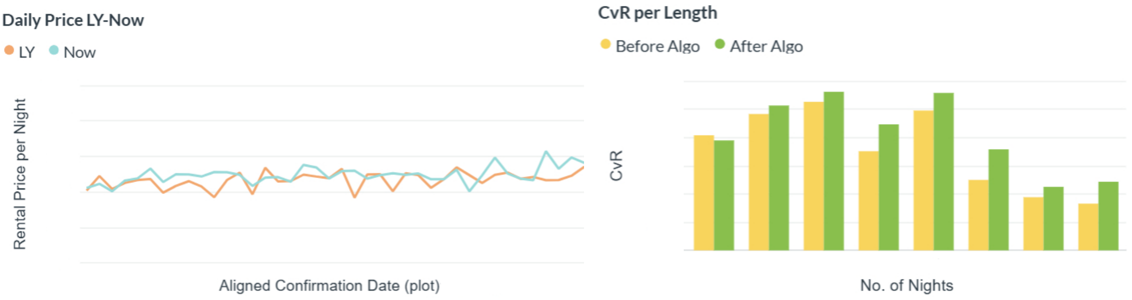

Figure 5 – Some Sections of the Daily Pricing per Length Dashboard (censored)

6.3. Indie Specials

Another dashboard was created for the relocation-deals known as Indie Specials (see Fig. 6). Their monitoring focused on the deals page funnel behaviour, since these offers behave differently from the routine rentals and remain sensitive to timing and the traffic quantity/quality. Filters were provided for the CI and CO dates, the CI and CO depots, the regions and the date entered; so that the questions could be answered through the various types of location and time window views.

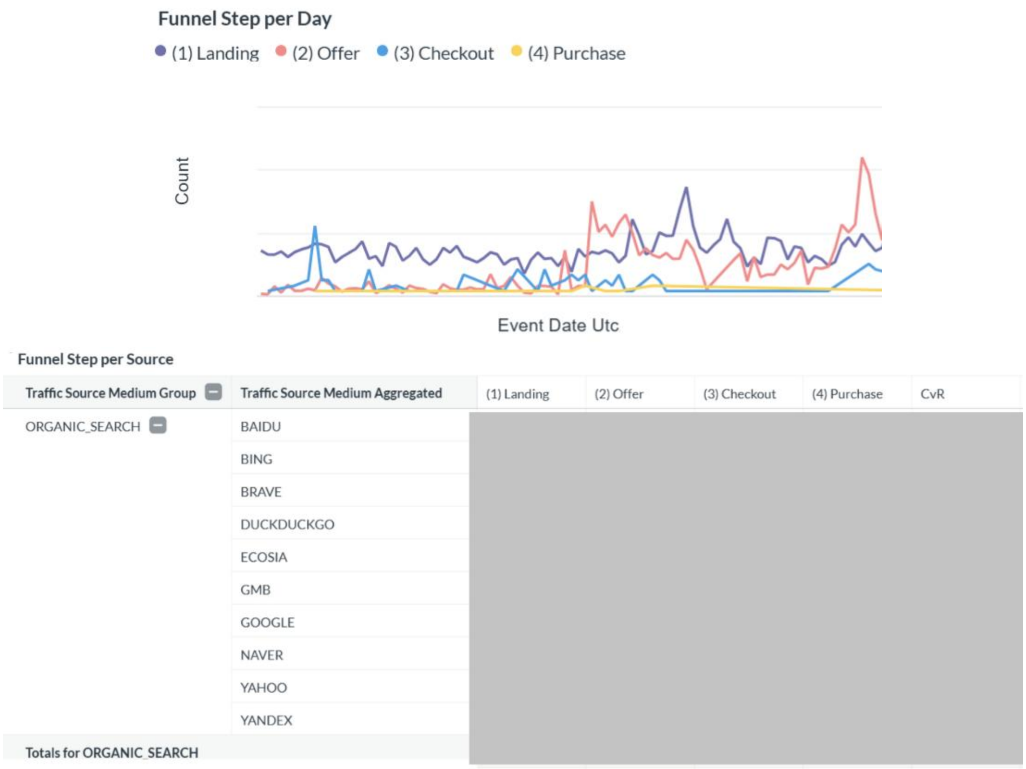

Figure 6 – Segments from the Indie Specials Dashboard (censored)

The main visual was a funnel-step time-series (landing, offer, checkout and purchase), because the first requirement was to see where the drop-offs occur once a wave of deals is released. The depot-level and the CI-CO location combination views were included because the relocation success is more directional than volume-driven. The deal status by route was also included, to show which combinations were open versus which were already booked, supporting the decisions on whether the deals should be injected again or redistributed. A meeting with the growth stakeholders was held to align on what a healthy source mix should look like -and as a result- a funnel-step volume by the traffic source (like social media advertisement or a search engine) table was also included; so that the campaign behaviour could be reviewed also by the quantity and the quality of the source of the traffic to the deals page.

7. CONCLUSION

7.1. Discussion

The internship made a valuable contribution to the practical development of the algorithmic tools and the dashboards that were aimed at enhancing the revenue decision-making process at Indie Campers. By integrating these systems into the daily operations, particularly in the fleet and pricing management, the objective was to improve the consistency and operational visibility for the company's central revenue team. The algorithms developed (which are the one-way pricing, discounts and vans-left tags and the search ranking algorithms) were created with the focus on both automation and interpretability, ensuring that the stakeholders had the necessary accessible and responsive tools to make informed decisions while maintaining the flexibility to make the adjustments when needed.

A central trade-off that emerged throughout the project was the balance between automation and the necessity for the human oversight. While the goal was to bring efficiency and repeatability to decision-making, it was clear that the tools should not replace human judgement entirely. Instead, the objective was to enhance the human decisions with the well-structured and data-driven insights. The dashboards became essential in providing the stakeholders with the visibility required to track the outcomes of the algorithms and to adjust the parameters in response to the shifting demand and the inventory conditions.

The development of these tools has facilitated the operationalisation of the theoretical concepts, like dynamic pricing and network optimisation, bringing them to life in a practical way. As such, the work represents an important step towards establishing a more structured, data-informed approach to the revenue management.

7.2. Limitations of the Study

Despite the comprehensive nature of the applied tools and methods, there were several limitations that affected the scope of the analysis. A primary constraint was the confidentiality surrounding internal operational data, which restricted the ability to share the detailed figures or fully assess the effectiveness of the tools beyond the company's scope. This limitation impacted the depth of most of the analyses.

Furthermore, the data quality and consistency issues arose due to the missing or the ambiguous historical data. For instance, the incomplete vehicle category groupings and the inconsistencies in the vehicle transmission data limited the ability to fully assess the impacts of the upgrades and downgrades. The lack of the standardised data governance across the company, further complicated the analysis.

Another limitation stemmed from the time constraints inherent in the internship period, which hindered the ability to evaluate the long-term effects of the tools implemented. While the algorithms and the dashboards provided the immediate operational benefits, their full impact could only be gauged after more time had passed and the tools had been fully integrated into the company's processes.

Additionally, the work was carried out within the specific context of the Indie Campers' business model, which may limit the broader applicability of the findings. While the developed tools are scalable and adaptable to the other locations and vehicle categories, the future applications will require ongoing adjustments to ensure they continue to meet the evolving business needs.

7.3. Recommendations

Looking forward, several recommendations can be made to build on the work completed during the internship and ensure its continued success. Firstly, the establishment of the clear governance structures for the algorithm deployment is essential. This will ensure that the changes made to the models are properly documented, tracked and aligned with the company's strategic goals. The regular audits and the reviews of the dashboards and the algorithms will be necessary to maintain their relevance as new data sources and business requirements emerge.

Moreover, a more robust data governance framework should be implemented to address the inconsistencies and the gaps identified during the projects. This would help standardising the data used across the company, improving the reliability of the outputs produced by the decision-support tools. With better data quality, the algorithms could be further refined to provide even more precise and actionable insights.

Although the algorithms and the manual inputs were designed with scalability in mind, it is recommended that the company continue to assess and to adapt these systems as the new locations and vehicle categories are added. A periodic review process should be established to evaluate the continued effectiveness of the models in the various operational contexts, ensuring they remain flexible and aligned with the shifting market conditions.

Finally, the company should consider investing in further automation to reduce the possibility of error by the manual interventions. Although the human-in-the-loop approach has been valuable in maintaining the control and flexibility, further automation could streamline the processes and make them more efficient. For example, replacing the spreadsheets with forms could reduce the likelihood of the human errors and optimise the workflow. Besides that, the enhancement of the user interface and the reporting capabilities of Metabase dashboards could provide the stakeholders with even more intuitive and actionable insights.

In summary, the work completed during the internship has significantly contributed to the Indie Campers' revenue management processes. However, ongoing efforts will be required to fully realise the potential of the implemented systems. By addressing the identified limitations and implementing the recommended improvements, the company can continue refining its approach to the fleet management, pricing and revenue optimisation.

References

Ampountolas, A. (2016). Revenue management, dynamic pricing and social media in the tourism industry: A case study of the name-your-own-price mechanism [Thesis, University of Exeter]. https://ore.exeter.ac.uk/articles/thesis/Revenue_Management_Dynamic_Pricing_and_Social_Media_in_the_Tourism_Industry_A_Case_Study_of_the_Name-Your-Own-Price_Mechanism/29723945/1

Bandalouski, A. M., Kovalyov, M. Y., Pesch, E., & Tarim, S. A. (2018). An overview of revenue management and dynamic pricing models in hotel business. RAIRO - Operations Research, 52(1), 119–141. https://doi.org/10.1051/ro/2018001

Bernardo, B. M. V., Mamede, H. S., Barroso, J. M. P., & Dos Santos, V. M. P. D. (2024). Data governance & quality management—Innovation and breakthroughs across different fields. Journal of Innovation & Knowledge, 9(4), 100598. https://doi.org/10.1016/j.jik.2024.100598

Chalupa, S., & Petricek, M. (2025). Application of revenue management practices in short-term rental management. Journal of Revenue and Pricing Management, 24(4), 326–333. https://doi.org/10.1057/s41272-024-00516-3

Guillen-Aguinaga, M., Aguinaga-Ontoso, E., Guillen-Aguinaga, L., Guillen-Grima, F., & Aguinaga-Ontoso, I. (2025). Data quality in the age of AI: A review of governance, ethics, and the FAIR principles. Data, 10(12), 201. https://doi.org/10.3390/data10120201

Hao, S., & Huang, L. (2023). How the time-scarcity feature of live-streaming e-commerce affects impulsive buying直播电商的时间性稀缺特征如何影响冲动购买. The Service Industries Journal, 43(11–12), 875–895. https://doi.org/10.1080/02642069.2023.2185231

Ishizaka, A., & Siraj, S. (2018). Are multi-criteria decision-making tools useful? An experimental comparative study of three methods. European Journal of Operational Research, 264(2), 462–471. https://doi.org/10.1016/j.ejor.2017.05.041

Khong, I., Aprila Yusuf, N., Nuriman, A., & Bayu Yadila, A. (2023). Exploring the impact of data quality on decision-making processes in information intensive organizations. APTISI Transactions on Management (ATM), 7(3), 253–260. https://doi.org/10.33050/atm.v7i3.2138

Kumar, S., Datta, S., Singh, V., Datta, D., Kumar Singh, S., & Sharma, R. (2024). Applications, challenges, and future directions of human-in-the-loop learning. IEEE Access, 12, 75735–75760. https://doi.org/10.1109/ACCESS.2024.3401547

Ladeira, W. J., Lim, W. M., De Oliveira Santini, F., Rasul, T., Perin, M. G., & Altinay, L. (2023). A meta-analysis on the effects of product scarcity. Psychology & Marketing, 40(7), 1267–1279. https://doi.org/10.1002/mar.21816

Li, D., Pang, Z., & Qian, L. (2023). Bid price controls for car rental network revenue management. Production and Operations Management, 32(1), 261–282. https://doi.org/10.1111/poms.13836

Nadj, M., Maedche, A., & Schieder, C. (2020). The effect of interactive analytical dashboard features on situation awareness and task performance. Decision Support Systems, 135, 113322. https://doi.org/10.1016/j.dss.2020.113322

Nunes, F., Alexandre, E., & Gaspar, P. D. (2024). Implementing key performance indicators and designing dashboard solutions in an automotive components company: A case study. Administrative Sciences, 14(8), 175. https://doi.org/10.3390/admsci14080175

Park, S., Bekemeier, B., Flaxman, A., & Schultz, M. (2022). Impact of data visualization on decision-making and its implications for public health practice: A systematic literature review. Informatics for Health and Social Care, 47(2), 175–193. https://doi.org/10.1080/17538157.2021.1982949

Sahoo, S. K., & Goswami, S. S. (2023). A comprehensive review of multiple criteria decision-making (MCDM) methods: Advancements, applications, and future directions. Decision Making Advances, 1(1), 25–48. https://doi.org/10.31181/dma1120237

Schwartz, J. M., George, M., Rossetti, S. C., Dykes, P. C., Minshall, S. R., Lucas, E., & Cato, K. D. (2022). Factors influencing clinician trust in predictive clinical decision support systems for in-hospital deterioration: Qualitative descriptive study. JMIR Human Factors, 9(2), e33960. https://doi.org/10.2196/33960

Sele, D., & Chugunova, M. (2024). Putting a human in the loop: Increasing uptake, but decreasing accuracy of automated decision-making. PLOS ONE, 19(2), e0298037. https://doi.org/10.1371/journal.pone.0298037

Smyth, C., Dennehy, D., Fosso Wamba, S., Scott, M., & Harfouche, A. (2024). Artificial intelligence and prescriptive analytics for supply chain resilience: A systematic literature review and research agenda. International Journal of Production Research, 62(23), 8537–8561. https://doi.org/10.1080/00207543.2024.2341415

Wissuchek, C., & Zschech, P. (2025). Prescriptive analytics systems revised: A systematic literature review from an information systems perspective. Information Systems and E-Business Management, 23(2), 279–353. https://doi.org/10.1007/s10257-024-00688-w

Zamfir, M. D. (2024). Scarcity effect and consumer decision biases: How urgency influences the perceived value of products. Journal of World Economy, 3(4), 27–34. https://doi.org/10.56397/JWE.2024.12.04

APPENDICES

Note: The appendices are written at a descriptive level to preserve the confidentiality.

A1. One-Way Pricing Algorithm Data Sources, Variables, and Data Preparation Steps

Dataset / Dataframe	**How it was acquired**	**Grain and typical volume**	**Key variables (high level)**
Availability & utilisation history (uti_data)	SQL extract from internal database into a dataframe inside the notebook	Daily grain, ~12-14 months horizon, typically a few hundred thousand rows	Region, country, CI depot, CO depot, vehicle category, date, utilisation, free vehicles
Historical one-way trip length (avg_length_data)	SQL extract of historical bookings, filtered to one-way rentals, then aggregated	Route grain (CI-CO), typically a few thousand routes	CI depot, CO depot, trip nights, booking counts, derived route average nights
Active depot-vehicle category combinations (open_combinations)	SQL extract of active depots and categories used as a master list	Depot x vehicle category, typically around one thousand rows	Region, country, depot, vehicle category, vehicle class
Distance and zoning maps (static files)	Read from controlled files stored in a shared drive	Small tables	Distance bands, zone, route validity
Min/Max fee boundaries and Average fee (min/max/avg sheets)	Read from controlled input sheets and used as hard bounds	Small matrices	Minimum, maximum and average OW fees by route
Manual override rules (OW manual input sheet)	Online sheet read during run, applied after fee computation	Typically tens-hundreds of rules	CI-CO pair, vehicle category, apply-from/until, manual price, conditional free-vehicle thresholds

Remarkable data preparation steps:

- Offer/test-like depots were excluded and only the true one-way rentals were kept (CI ≠ CO), with cancelled bookings removed from the trip-history extract.
- Route-night history was cleaned by filtering the irregular sources/substitution artefacts, then protected against the small samples by using a zone-level proxy when the route history was weak.

- Utilisation extracts were restricted to the latest snapshot and a fixed forward horizon, with the negative/invalid fleet states filtered out.
- Key fields were normalised (case/whitespace) before merges and the mapping tables were de-duplicated to prevent accidental row multiplication.
- Zone-crossing routes were removed for the route-length layer to keep the route averages comparable.
- Delta-type adjustment logic was bounded and the missing utilisation points were handled through explicit fallback branches, so the sparse data could not create extreme fees.

A2. Discounts & Vans-Left Tags Algorithm Data Sources, Variables, and Data Preparation Steps

Dataset / Dataframe	**How it was acquired**	**Grain and typical volume**	**Key variables (high level)**
Availability & utilisation lookahead (uti_data)	SQL extract into dataframe inside the notebook	Daily grain, ~12 months lookahead, typically ~300k rows	Region, country, depot, vehicle category, date, booked, available, free, utilisation
Weekly aggregation features	Derived in-notebook by grouping dates into week buckets	Depot x vehicle category x week, 52-week horizon	Weekly free vehicles, weekly utilisation, vehicle category counts
Strategy matrices (discount/tag rules)	Read from controlled files or notebook constants	Small tables	Week-bucket logic, vehicle category counts, utilisation quartiles, discount/tag allocations
Manual override rules (discount/tag sheet)	Online sheet read during run, applied after assignment	Typically tens of rules	Depot-vehicle category scope, apply-from/until, discount level, vans-left tags

Remarkable data preparation steps:

- Temporary depots and the offer/test depots were filtered out in the SQL query, only the latest snapshot and a fixed lookahead horizon were used.
- Empty/null vehicle category identifiers were removed and the category-class mapping used an explicit fallback label so the mapping gaps stayed visible without breaking the run.

- Weekly roll-ups were rounded before the rule-table mapping, so the small floating noise did not change the discount/tag decisions.
- Vehicle category-count logic excluded the "zero free" categories, avoiding the artificial inflation of the depot variety.
- Manual override sheets were cleaned with the missing mandatory-field drops, preventing the open-ended or partially filled edits from silently overriding the outputs.
- Deduplication was applied before the key merges to avoid duplicated joins expanding the output.

A3. Search Ranking Algorithm Data Sources, Variables, and Data Preparation Steps

Dataset / Dataframe	**How it was acquired**	**Grain and typical volume**	**Key variables (high level)**
Forward utilisation & availability (uti_data)	SQL extract into dataframe inside the notebook	Daily grain; ~12 months horizon, typically ~330k rows	Region, country, depot, vehicle category, vehicle class, date, utilisation, free vehicles
Recent booking history (booked_totals)	SQL extract of recent bookings, aggregated into rolling windows	Depot x vehicle class/vehicle category x window, typically a few thousand rows	Booking counts by period, DE, CI depot, vehicle category
Active combinations master list (open_combinations)	SQL extract used as output grid	Depot x vehicle category, typically ~1k rows	Active depot-vehicle category list, vehicle class mappings
Manual override rules (ranking sheet)	Online sheet read during run, applied at the end	Typically a small set of rows	Scope (region/country/depot/vehicle category), apply-from/until, rank, conditional free-vehicle thresholds

Remarkable data preparation steps:

- Booking history extracts excluded the offer/test depots and the cancellations, used a fixed recent window, and filtered out the irregular/substitution artefacts so the "preference" signal stayed clean.
- Utilisation extracts excluded the temporary depots and were restricted to the latest snapshot + forward horizon to prevent mixed states.

- Missing booked/utilisation values were filled early, and the rank fallbacks were applied within groups so the missing data did not jump to extreme ranks.
- A complete output grid was enforced from the active-combinations master list so all depot-vehicle category combinations were included just in case.
- Export columns were dtype-stabilised (integers for scores, defaults for blanks) to prevent the downstream import issues.
- Manual override rows were cleaned and applied with the specificity rules (depot > country > region), reducing the accidental overrides.

A4. Minimum Nights Dashboard

Minimum bookable number of nights rules -that were designed to decrease the amount of gap days in high season in order to increase revenue- monitoring was built by using the SQL-based question and visualisation creation in Metabase (see Fig. 7). It was created mainly for the CPT associate usage, since the minimum nights settings were adjusted manually and their visibility was required to avoid any inconsistent or wrong application across the network. The filters were designed around the operational questions, consisting of: Region, country, depot, vehicle class, PU month and PU week.

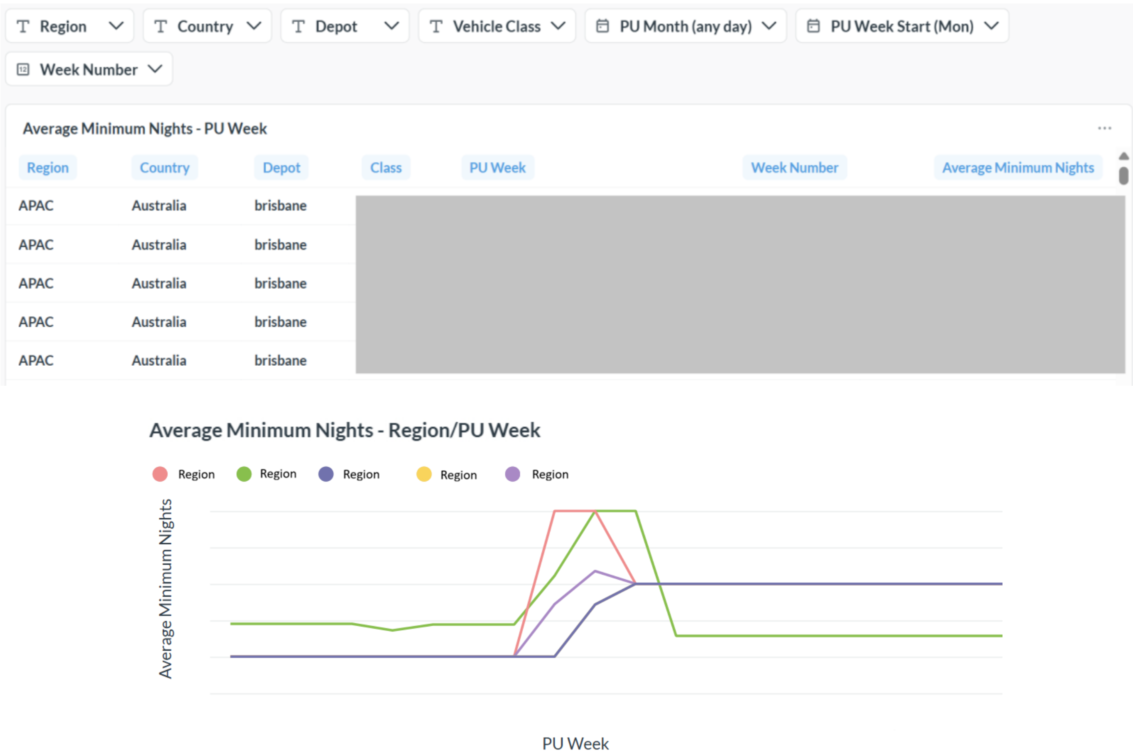

Figure 7 – Items from the Minimum Nights Dashboard along with the Filters (censored)

A “what is live” table (average minimum nights by the PU week, depot and vehicle class) was paired with the trend views showing the average minimum nights by the PU month and PU week, beginning at the region level and then drilling down to country and depot levels. The weekly aggregation was used to reduce the noise and to match how these decisions were usually reviewed. Readability and platform-related constraints emerged when too many depot series were tried to be displayed, so some parts of the dashboard were structured to rely on scoping through filters. The KPI choices and the plot priorities were aligned in a short review with the CPT associates, based on what they want to see first when the minimum nights appear too strict or too loose.