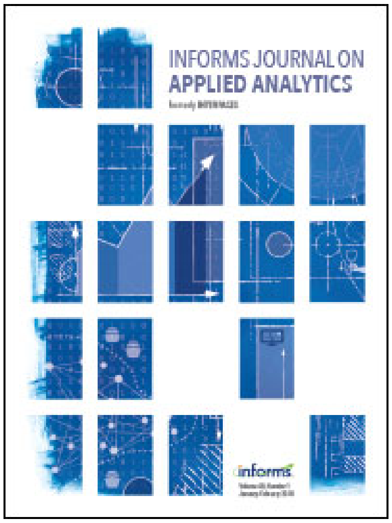

INFORMS Journal on Applied Analytics

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Carnival Optimizes Revenue and Inventory Across Heterogenous Cruise Line Brands

Justin Beck, John Harvey, Kristina Kaylen, Corrado Sala, Melinda Urban, Peter Vermeulen, Norman Wilken, Wei Xie, Dan Iliescu, Pratik Mital

informs
http://pubsonline.informs.org/journal/inte

INFORMS JOURNAL ON APPLIED ANALYTICS
Vol. 51, No. 1, January–February 2021, pp. 26–41
ISSN 0092-2102 (print), ISSN 1526-551X (online)

THE FRANZ EDELMAN AWARD
Achievement in Operations Research

Carnival Optimizes Revenue and Inventory Across Heterogenous Cruise Line Brands

Justin Beck,[a] John Harvey,[a] Kristina Kaylen,[a] Corrado Sala,[a] Melinda Urban,[a] Peter Vermeulen,[a] Norman Wilken,[a] Wei Xie,[a] Dan Iliescu,[b] Pratik Mital[b]

[a] Carnival Corporation & plc, Miami, Florida 33178; [b] Revenue Analytics Inc., Atlanta, Georgia 30339

Contact: justin.beck@pocruises.com.au (JB); jharvey@hollandamerica.com, https://orcid.org/0000-0002-3218-2190 (JH); kkaylen@carnival.com (KK); csala@carnival.com (CS); melinda.urban@carnivalaustralia.com (MU); peter.vermeulen@carnivalukgroup.com (PV); norman.wilken@aida.de (NW); wxie@princesscruises.com (WX); diliescu@revenueanalytics.com (DI); pmital@revenueanalytics.com (PM)

Received: October 5, 2020
Accepted: October 5, 2020
https://doi.org/10.1287/inte.2020.1062

Abstract. Carnival Corporation & plc identified the need for a cutting-edge revenue management system; however, existing solutions from the airline and hospitality industries were not compatible with the idiosyncrasies of the cruise domain. As such, the company partnered with revenue analytics to build a complete revenue and inventory management system to meet its requirements. Yield optimization and demand analytics (YODA) is a system that leverages a unique quadratic programming model to jointly determine cruise prices and allocate cabin inventory to multiple cruises (e.g., 14-day and 7-day lengths) offered simultaneously on a given ship. The optimization inputs come from several machine learning algorithms that predict demand. YODA combines these algorithms with an elasticity model derived from an exponential curve to represent the unique price-sensitivity behavior observed in the cruise industry. The system generates millions of price recommendations each day and has been used to price voyages on 65 Carnival ships, approximately one quarter of the ships in the entire cruise industry, since December 2017. During A/B testing, YODA generated a 1.5%–2.5% incremental uplift in net ticket revenue, which is a significant revenue increase because Carnival was a Fortune 300 company in 2019.

Keywords: revenue management • cruise lines • machine learning • optimization • quadratic programming • CART • Edelman Award

Introduction

Overview of the Industry and Carnival Corporation

In 2019, more than 30 million people, including 13 million from North America, 6 million from Europe, and 4 million from Australasia, cruised on more than 270 ships (Cruise Lines International Association 2019). Of those passengers, nearly 50% sailed on Carnival Corporation ships.

Carnival Corporation is the world's largest leisure travel company, providing guests worldwide with extraordinary holiday experiences at exceptional value. With operations in North America, Australia, Europe, and Asia, its portfolio features Carnival Cruise Line, Princess, Holland America Line, Seabourn, P&O Cruises (Australia), Costa, AIDA, P&O Cruises (UK), and Cunard. Together, the corporation's cruise lines operate 105 ships visiting more than 700 ports around the world. Each year, more than 265,000 daily cruise guests spend approximately 93 million days sailing aboard the Carnival fleet, which includes a staff of 100,000 shipboard employees. Throughout the paper, we use the terms *voyages* and *cruises* interchangeably.

The first six Carnival brands listed in Table 1 joined the corporation through acquisitions at different points in time; before yield optimization and demand analytics (YODA's) implementation, each operated largely independently and differed significantly on all operational aspects from information technology (IT) systems to organizational structures. In particular, they had separate revenue management (RM) teams that reported up to separate chief commercial officers. Project collaboration was limited to occasional conferences. The Carnival Corporation had never previously attempted a major cross-brand project.

Revenue Management in the Cruise Industry

Revenue management models and techniques were first developed and used by American Airlines (Belobaba 1987, Smith et al. 1992) to help it address increased competition and changing market conditions, after the United States enacted its Airline Deregulation Act in 1978, through a combination of purchase restrictions and capacity-control fares (Talluri and van Ryzin 2004).

Table 1. Each of the Six Carnival Brands that Were Part of the Original YODA Project Represents Different Markets and Has Unique Business Challenges

Brand (year originated)	Markets	Ships (capacity)
Princess Cruises (1965)	Serves a global market on ships that sail throughout the world, including North America, Europe, Asia, and Australia	20 ships (686–3,700 guests)
Holland America Line (1873)	Serves largely North American guests, who are looking to explore and adventure, particularly specialist tours and experiences in Alaska	15 ships (1,260–2,670 guests)
P&O Cruises UK (1837)	Serves the UK market on ships that sail throughout the world, with many air-inclusive products	8 ships (1,880–5,200 guests)
P&O Australia (1932)	Serves the Australian and New Zealand market on ships that sail throughout Oceania	6 ships (680–1,450 guests)
Cunard (1840)	Primarily serves UK and U.S. markets on ships sailing throughout the world	3 ships (2,000–2,700 guests)
Seabourn (1987)	Serves ultra-luxury guests on small suite-only ships that sail to unique locations and provide specialized itineraries	6 ships (270–600 guests)
AIDA Cruises (1996)	Serves primarily the German market on ships sailing throughout the world	14 ships (1,300–6,600 guests)

Note. AIDA adopted the YODA system in 2019 after the system was deployed successfully for the first six brands.

The RM concept was to optimize the setting of prices over time to maximize revenue (yield) from selling perishable services (e.g., seats on a flight). The success of RM in the airline industry was replicated in hotels with Marriott International's implementation of fenced rate discounts (i.e., discounted rates for meeting specified conditions) and availability controls, leading to widespread adoption of yield strategies across all major North American hotel chains (Cross et al. 2010). Applications such as Holiday Retirement's use of RM to determine rents for senior living communities (Kuyumcu et al. 2018), airline inventory control across a variety of fare structures (Fiig et al. 2009), Marriot International's competitive bid-response pricing (Hormby et al. 2010), and Europcar's setting of prices for vehicle rentals (Guillen et al. 2019) enhanced traditional RM models with customer-centric pricing models and techniques that allowed RM proliferation to cruise lines, areas of travel such as rail and rental cars, and entertainment.

Although the cruise industry is one of the most dynamic parts of the travel and hospitality sector, it did not receive much attention in the field of revenue management until recently (Ayvaz-Cavdaroglu et al. 2019). Even then, it was not in the context of jointly allocating inventory and determining price, a unique facet of YODA.

RM in the cruise industry is similar to that of other industries but faces unique challenges. Similar to hotel room types, Carnival offers many cruise cabin categories at different price points. To avoid pricing volatility because of the small numbers of cabins in some categories, similar cabin categories on a ship are grouped based on their features; these include metas, such as inside cabins, oceanview cabins (window), verandahs (cabins with private sea-view balconies), and suites (luxurious cabins with larger private balconies). Booking-curve similarity analysis is used to split metas into submetas, where specific cabin categories exhibit different behaviors. For example, the oceanview meta may be split into obstructed oceanviews (e.g., the cabin has a window, but a lifeboat blocks the view) and standard oceanviews. Other similarities include determining the right level of promotional investment, handling overbooking and upgrades (similar to airlines), guest sourcing and market mix, and airline sourcing (similar to package holiday operators in Europe).

However, cruise revenue management has additional considerations. Cruising has a long booking cycle, where bookings can be taken more than two years before sailing, and last-minute purchases can occur the week of departure. Prices can fluctuate up and down throughout this booking horizon. Incentives are often offered exclusively to guests who book early, but some still wish to hold out for the possibility of a lower fare. This behavior is driven by historically high-to-low prices (i.e., large fare differences) in the cruise sector. However, these fare differences can over time result in unsatisfied guests; the price paid for the cruise is a frequent dinner conversation topic on cruises. To mitigate against this risk and that of cancellations, cruise lines also allow guests to book cruises using an initial deposit, with additional and/or final payments due two to three months before sailing, to encourage guests to book early at higher prices; however, a number of bookings will be cancelled for various reasons before sailing.

Another unique aspect of cruising is the ability to offer a variety of destinations and experiences bundled into one holiday package. This includes experiences at ports, spa treatments, casinos, a myriad of different dining options and bars, live music, and entertainment. This onboard revenue can represent up to 30% of a brand's overall revenue, highlighting the strong positive correlation between revenue optimality and occupancy. However, this is complicated by the unique occupancy constraints that exist in the cruise industry; ship capacity is constrained by both the number of cabins and the total number of passengers on the ship, because many cabins can be occupied by up to four guests.

On the ticket RM side, which is YODA's primary focus, selling the right product (i.e., voyage ticket) to the right guest at the right time for the right price (Cross 1997) is as important in cruising as it is in the RM function in any other industry. However, many cruise itineraries are unique, and multiple itineraries may be booked on a ship at any one time; for example, some passengers may be on a ship for 7 days and others may stay on for 14 days. These *networked* voyages may have different demand curves and price points. This requires precise identification of the relevant history from previous cruises, careful price management, and the implementation of appropriate duration controls to ensure that the allocated capacity on each subset of the cruise network aligns with contracted seats on flights to transport guests to the ship. For example, a guest who was looking for a two- to three-week cruise holiday to the Mediterranean during the summer of 2019 would have had 32 cruise choices sailing from Southampton on P&O Cruises UK alone, and would have been able to visit 19 ports.

When a ship supports multiple cruises on a same day of travel, it is necessary to allocate that ship's cabin inventory to those multiple cruises. The inventory allocation is managed at a submeta level of detail. These inventory allocations are then used to support planning; for example, for the UK brands, allocation might include the number of air seats required to transport guests to/from each cruise, as well as cabins in the reservation system, so Carnival can achieve the optimal demand across each of those cruises.

Unlike airlines and hotels where inventory is protected for high-yield business customers, last-minute bookings in the cruise industry often require deep discounting. Additionally, overbooking practices pose a much higher risk. Airlines frequently use point-of-departure buy-offs (i.e., incentives as compensation for booking alternative flights/travel) because passenger flights are shifted only a few hours. For cruises, a buy-off may require a passenger to wait weeks for a comparable offering.

Revenue Management Practice at Carnival Before YODA

Before the implementation of YODA, each Carnival brand managed revenue and made RM decisions using its own tools and processes. Although similar tools existed across the brands, they varied significantly from simplified forecasting and cancellation predictions to decisions made by analysts who used Excel tools to compare a cruise with a similar cruise based on historical data. These tools were capable of detecting large, sudden deviances from expected behavior. They were less adept at detecting slowly developing patterns or handling any quick incremental adjustments needed to maximize revenue.

A typical voyage may be on sale in two core markets and additional secondary markets, with four to seven submeta prices. The submetas could include 15 to 40 categories, which have varying locations and/or sizes. Prices also need to be managed for bookings in which different numbers of guests occupy a cabin. One voyage could require managing more than 2,000 different price points. This complexity meant that analysts had to choose where to focus their attention. Analysts had to limit regular monitoring to voyages close to departure and only in core markets and expend significantly less effort on voyages further from departure. Voyages would often show a significant difference in the prices at the time a cruise is first made available for booking compared with the prices offered during the last few weeks before sailing. Early analysis in the project showed that this approach did not result in maximizing revenue. To add to the complexities, analysts were responsible for publishing prices, a manual process that included a number of time-consuming quality checks. Inventory allocations would often be subjectively determined across overlapping cruises of different durations (e.g., 7 versus 14 nights), and these were rarely reviewed.

To address the strategic requirements and bring both consistency and science to RM across these brands, while handling all the previously listed complexities, the concept of a Carnival RM system was conceived. The YODA system would enable each brand to manage its own nuances while maintaining a consistent approach with its sister brands within the corporation. Because a viable off-the-shelf solution did not exist for the cruise industry, Carnival Corporation partnered with Revenue Analytics, a software and consulting company, to design and build its own system.

Brand Nuances

Each brand has unique aspects because each sells to different markets with distinct legal constraints and with customers who have unique needs and wants. Additionally, past behavior and the presence of multiple cruise lines create different market conditions, as we list here.

- P&O Cruises UK: Offers adult-only ships and fly-cruise packages, where the brand prepurchases airline seats and includes them in the overall price.
- P&O Australia: Frequently reaches the safety limit for total passengers so the upper-berth capacity must be carefully managed; its passengers typically include a number of families with children. In the remainder of this paper, lower berth refers to the first two beds in a cabin; upper berth refers to a bed other than a lower berth.

• Seabourn: Provides ultra-luxury experiences because it is a luxury brand with smaller ships that contain only suite cabins, which can frequently command much higher prices.

• Holland America and Princess Cruises: Provide a unique cruise and land experience in Alaska; their pricing among all their cruise and land options must be consistent.

• AIDA: Cannot raise prices above rates published in brochures because of German laws.

• Cunard World Cruises**:** Full-world cruises visit a number of ports over 100 or more days; these world cruises are sold as a package but can also be broken down into a number of sellable sectors, creating a complex network inventory-allocation problem.

In addition, the published prices in various markets and in multiple currencies differ because of regulations about including taxes, fees, and port expenses in the published price.

Developing Common Terminology and Rules

The YODA project was conducted by a cross-brand team of analysts, business leaders, and data scientists across Europe, Australia, and the United States, who were dedicated full time to the project and had defined roles and responsibilities. These individuals became the champions of YODA when they returned to their roles within their particular brands. Geographical diversity was one of the biggest barriers they had to overcome initially. Time zones are barriers to working across brands within Carnival. There was no hour of the day when the entire team was working concurrently.

Another challenge was reaching consensus across the brands because each brand's history and organizational structure differed. Overall, the team identified 80 categories of different business processes and data definitions to resolve before it could design and build a single system that Carnival could use consistently across all brands. These categories and ultimate decisions were all consolidated into a *brand alignment matrix* that showed how each brand's process and definition differed from those of the other brands and highlighted key elements that each brand required, so that no brand lost any key preexisting capability. The project team representing each brand worked closely with the teams from the other brands and with its respective brand's RM, finance, operations, legal, and IT teams to build consensus on items such as the following:

• Geographical classification: For example, should a cruise ship that sails from Sydney to Hawaii to Los Angeles and ends at Seattle be classified as an Australian, Hawaiian, Pacific, or U.S. West Coast cruise?

• The first day of the week: Should it be Sunday or Monday?

• The definition of *net* price: Should revenue generated be based on the fare paid minus expenses and the cost of add-ons?

• The legal requirements in different jurisdictions: For example, Germany and Japan have strict consumer protection laws, and many countries require the inclusion of taxes and fees as part of the advertised prices.

This process took two to three months, and the changes required ranged from data definitions to website layouts.

Design, Development, and Deployment of YODA

After converging on these issues, the YODA project planning began. It was divided into three phases.

Phase 1: Initial Exploratory Phase, Design, and Proof of Concept (August 2015 to November 2015)

The objectives of phase 1 were to solidify the understanding of the technology, data, analytics, and pricing architecture by brand and region and to converge on a vision of a solution and a roadmap to address all brand requirements. In addition, the team identified opportunities to quickly implement some changes (i.e., quick hits) to generate initial revenue gains. In a project planning workshop, team members built a detailed plan and took part in a data discovery exercise in which, for each brand, they assessed the data and systems that were readily available and identified the functions that needed to be built.

To identify quick hits, the team built an exploratory model to provide insights into process and strategy changes that would drive immediate revenue improvement. This model unconstrained booking curves (i.e., the cumulative materialization of when bookings have historically been made over time) to impute latent demand. This unconstraining was done by looking at similar historical cruises and identifying the cruises for which inventory had been sold out (i.e., there was unmet demand), or there was little booking activity in certain periods because inventory had been priced too high.

The team ran an optimization model on these unconstrained booking curves to understand which demand the optimization would choose, retrospectively, to maximize revenue. The resulting insights ranged from reducing overbooking on lower, less-expensive categories, to the revenue-optimal number of bookings allowed from different markets. The initial insights derived from identifying the quick hits allowed the YODA system to deliver an incremental revenue uplift of approximately 0.3%.

One of the key opportunities identified in the exploratory phase was the introduction of multiple fares available to the public at the same time. By offering a second, publicly available, discounted fare with fewer inclusions (e.g., premium beverages, Wi-Fi) and stricter terms and conditions compared with the preexisting (i.e., standard) fare, demand could be captured from guests with different personal preferences. This is a win-win situation because customers can choose the options that best suit their needs, and Carnival can maintain its revenue (Figure 1).

Phase 2: Build and Market Test Prototype Solution (November 2015 to September 2016)

The objectives of phase 2 were threefold:

1. Develop, test, and deploy the first version of YODA, including price and inventory optimization.
2. Gather feedback and identify analytics, user interface (UI), and business changes required for successful enterprise development and implementation.
3. Finalize enterprise YODA design and development plans, deploy the system, undertake user training, and develop associated change management processes.

Interactive development through weekly meetings and peer reviews ensured designs were not only scientifically sound but also validated them against real-world user expectations and experiences. This development took place in a series of iterations, which were focused on collaboratively designing and developing a subset of capabilities (demand forecast, market response models, price and inventory optimization, and UI), as shown in the arrows in Figure 2.

Approximately 15 data scientists from the six brands and Revenue Analytics, and approximately 20 data experts, business leaders, and project managers, regularly flew around the world to meet, review progress, and resolve any issues. Physically being together allowed them to establish the purpose of each part of YODA, the variables it should consider, and possible approaches for satisfying the requirements. Where multiple valid options existed, each was evaluated on its merits and by using prototypes, and the model that worked best was selected.

A pilot study and a Tableau UI were deployed to enable users to test a prototype system. At this stage, YODA was generating price recommendations for all voyages nightly. Analysts provided overwhelmingly positive feedback. Most importantly, this pilot helped the team to understand which areas were most important in driving price recommendations and which areas needed improvement.

This period was structured to facilitate A/B testing. Brands typically entered around 30%–40% of their inventory into the prototype to allow YODA to manage their prices, and each brand carefully selected its included cruises to fulfil two criteria. They must include the following:

1. A representative sample of the overall product that the brand offers with respect to different geographical groupings of cruises, a mix of air-inclusive and cruise products versus cruise only, and a representative mix of markets, ship types, and durations; and
2. A control group sample of *equivalent* products that would not be managed using YODA for the duration of the market test.

Although the team made every effort to conduct a pure A/B test, the realities of developing a real-world system presented complications. For example, if a brand wanted to test YODA's ability to manage the pricing of Caribbean cruises, and two ships were positioned in the area for a specific period, an analyst could manage one ship using YODA and another analyst could manage the second ship using legacy tools. However, these analysts were members of the same RM team and may have attended many of the same meetings at which pricing was discussed. Thus, they could have been influenced by these discussions. As soon as the analysts were confident that YODA was making better decisions, they sought to have it manage all their products as soon as possible. Although this was a great vote of confidence for YODA, it invalidated the A/B tests in some circumstances.

Figure 1. (Color online) Optimization by Multiple Segments Allows Guests to Make Trade-offs and Purchase Cruise Packages at the Price Point They Wish to Pay

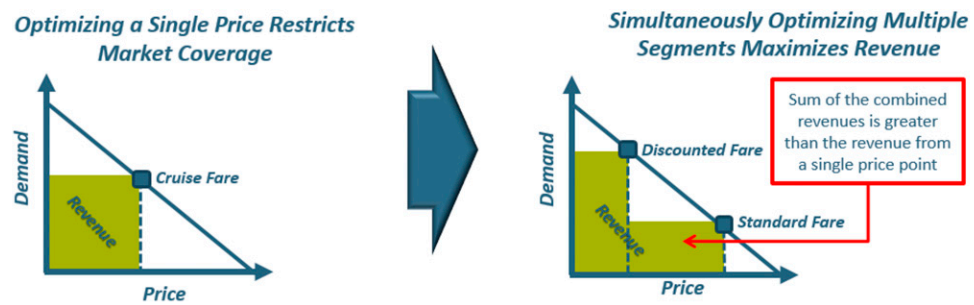

Note. The area under the curve (revenue) is greater when the two price points are available.

Figure 2. (Color online) Phase 2 Project Plan Was Divided into Four Phases with Different Lead Areas of Focus, Each Comprising a Number of Two-Week Development Iterations with Different Areas of Focus, Plus a Testing Phase and a Market Test Phase

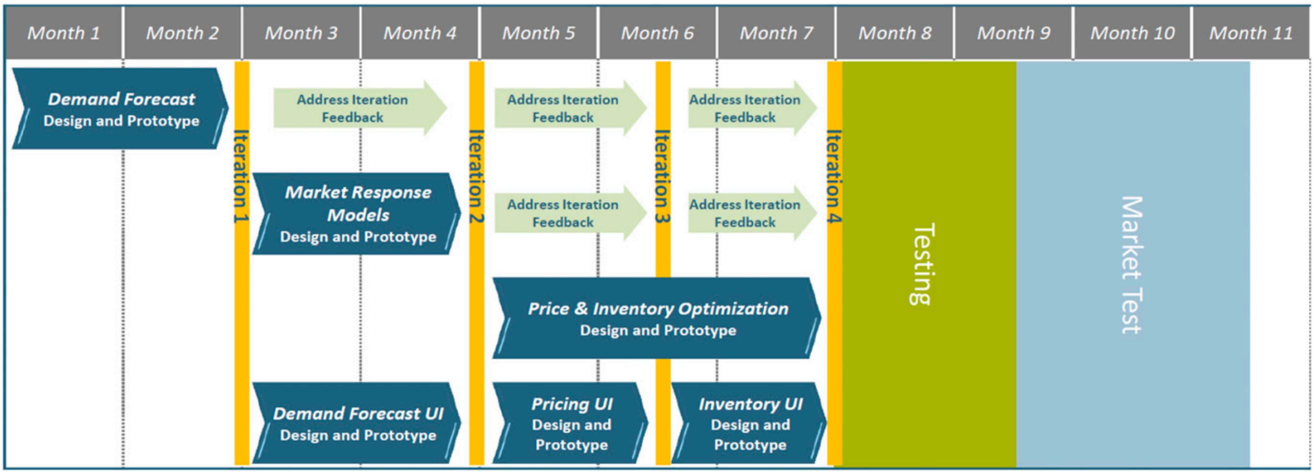

The selection was harder for smaller brands. For example, Cunard has only three ships, and they all tend to offer different itineraries. Therefore, Cunard required a solution tailored to its needs. The team decided to develop six alternative A/B testing systems to measure uplift over a three-month period and designed them such that they handled the nuances of the different brands while adjusting for confounding variables that could not be excluded.

In addition to the A/B testing results, several measures were used to determine the success of the YODA prototype. These included use and adoption from the analysts and key performance indicators (KPIs) relating to the component modules. The KPIs for several modules, including demand and retention forecast models, used standard backcasting approaches to evaluate accuracy using weighted mean average percentage error (WMAPE). Backcasting approaches included running 2019 cruises through YODA using data available as of December 2018 and then comparing YODA's expectations with actual occurrences in 2019. This involved running the models with all data as of one year earlier, predicting the events of the next year, and then comparing these results with the actual results that materialized in that period.

Phase 3: Enterprise Solution, Design, Beta Testing, and Rollout (October 2016 to May 2018)

Although some peripheral components were developed during phase 3, the focus in this phase moved to fine-tuning and building configurations. The project team monitored the impact these changes had on the KPIs and then turned its efforts to dealing with special cases. In particular, a key component of this phase was the parameterization of the models within YODA. YODA is highly configurable, with thousands of parameters available to guide the analytics toward specific outcomes based on business knowledge. We conducted a market test to find the parameters that worked best for the overall set of products being tested. In this phase, parameters were fine-tuned, especially in situations where data scarcity prevented the effective use of data-intensive machine learning techniques.

Improving these parameter settings and determining where they should be applied became an iterative process that looked at the worst-performing voyages or trades (e.g., voyages in a similar location such as Alaska or the Caribbean) in a given KPI, investigating its key drivers, and solving for the best choice of parameters. With the backcasting approach used in forecasting, this exercise was straightforward to execute.

Other changes and opportunities for refinement resulted from a science health check report (HCR). The HCR gave additional confidence to scientists and analysts that a YODA recommendation was sound via a series of self-diagnostic tests. Each recommendation received an automated quality classification of high, medium, or low, depending on its reliability and the confidence the scientists and analysts had in the recommendation. For example, because of the changing nature of cruises (e.g., a ship might be new or it could be in a different geographic location every year), some voyages have a less reliable history and will always need more oversight by analysts. As another example, if the recommended price changes significantly due to either elasticity values or capacity constraints, the change requires an analyst review prior to being published online.

The HCR looks at each of YODA's separate model inputs and outputs for business reasonability (e.g., history of reliability, price changes within specified bounds, forecasts within specified percentages of capacity) and evaluates their impacts on the overall recommendations. These checks are extensive and range from having enough historical voyages to having a forecast for each submeta and market to having business-reasonable elasticities. The HCR provided a useful diagnostic tool to support scientific testing and eventually became so popular that it was added to the UI and automated within YODA.

The team placed a significant focus on the UI development. A number of interactive design sessions took place involving experts from the analyst community and the development team to ensure that the interface satisfied the needs and wants of all the end users. The YODA design also included a scenario modeler capability—a UI that allows users to evaluate the effects of input changes on YODA's recommendations. YODA's core engine and UI features went into production for the six YODA brands in December 2017, and all six brands have since used it continually. Additional features of the UI were designed and deployed up to May 2018.

Significant design features of the UI include ease of use for multiple users in different countries, time zones, and business organizations; compliance with regulations; clarification of assumptions and inputs to the optimization model; and the ability to efficiently send price updates to POLAR, the YODA brands reservation system, thus saving thousands of people-hours annually.

After the implementation for the six brands was complete, Carnival Corporation decided to bring AIDA, Carnival's fastest-growing brand, onto YODA (Table 1). AIDA brought unique problems to solve because its cruises can be sold to the German market either as cruise-only voyages or as packages bundled with air fares. AIDA's 14 ships also frequently carry large numbers of children in upper berths. Testing AIDA validated YODA's transportability to other cruise lines, and within five months, in December 2019, AIDA was successfully integrated into YODA. The integration activities included loading of any additional data sources, transformation of data into existing staging tables that are data-source agnostic, and the configuration of existing analytical modules to handle AIDA's unique business needs.

Data and Infrastructure

The flow diagram in Figure 3 summarizes YODA's data infrastructure. POLAR is its primary data source; however, the YODA data warehouse (DWH) also takes supplementary data feeds from other data warehouses

Figure 3. (Color online) YODA's Infrastructure Combines a Wide Variety of Data Sources into a Single Source of Truth for Users

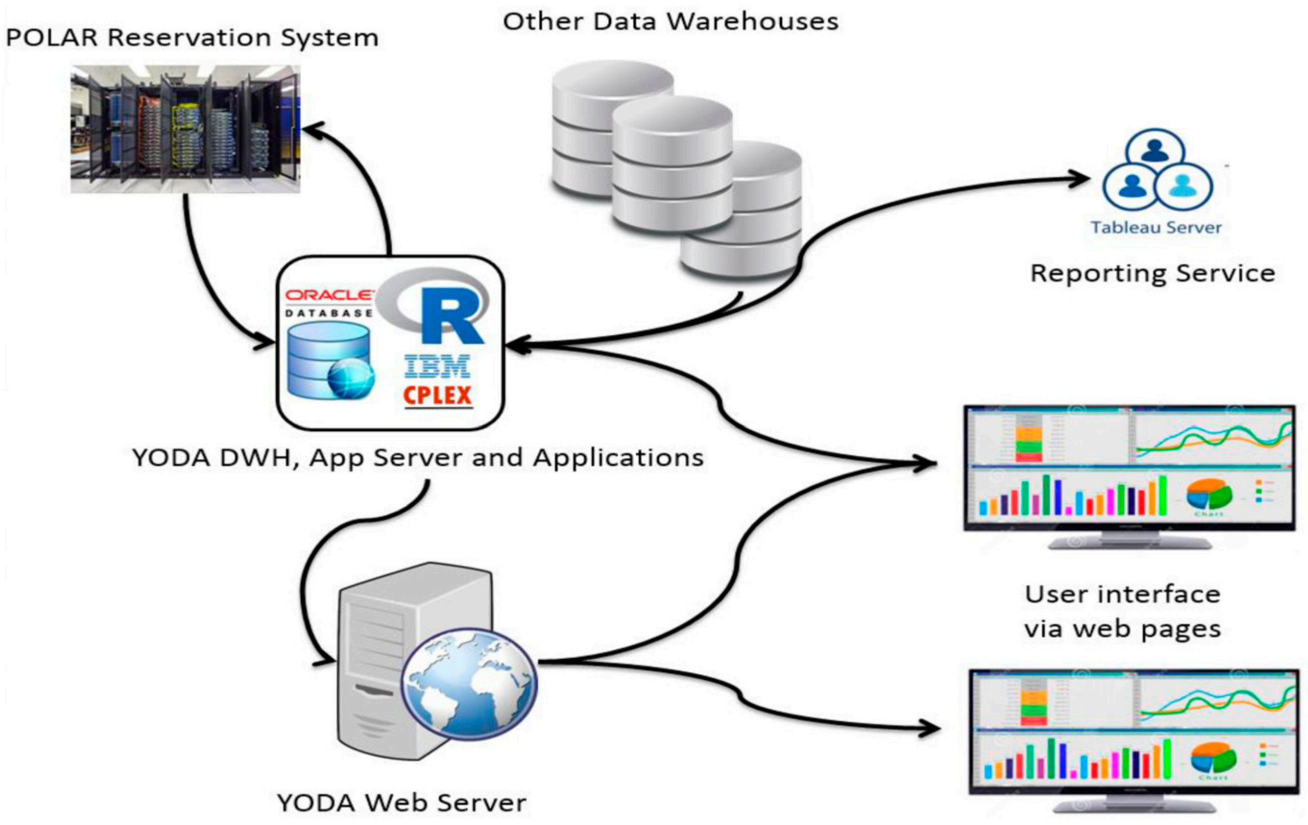

that include data not held in the reservations system. The data feeds coincide with the end of the business day in the home locations of each brand that uses YODA. For each brand, the batch process runs during that brand's night hours; this process includes updating the UI and generating reports containing the most current information by the start of the following business day.

YODA runs different models with varying frequencies. Macro-level models (e.g., segmentations and price-elasticity derivations) execute quarterly. In contrast, models that change from week to week (e.g., retention forecasts, developing trends, price normalization, and booking-curve updates) run weekly. Updating the capacity with the latest booking position and generating updated price recommendations and inventory controls is done nightly. This reduces the run time of the daily batch process and brings stability to some of the components. YODA is comprised of 29 modules, containing 165,000 lines of code, which generate more than five million price recommendations daily. These modules process 111 billion rows of data contained in 2,200 tables, including source data and transformation tables. The source data are cleaned and checked and merged with other source and reference data sets to create a final set of data ready to be integrated, or staged, into YODA. These final data sets are loaded into the staging tables in YODA and are collected at the start of each batch. These recommendations are accessed routinely and reviewed by more than 200 users across the business and IT teams, in addition to the 15 data scientists across the six brands' operations research teams.

At the core of YODA are the database servers and application servers. Database servers process the raw data into inputs that the application servers need and store user changes and price recommendations to be submitted to POLAR. Application servers consist of R programs, CPLEX solvers, Python programs, and shell scripts to execute the models. The UI layer contains the YODA website and reporting portal based on Tableau.

The YODA architecture is source-system agnostic and can handle multiple reservation systems and data structures within the cruise industry. This allows the integration of other cruise lines into existing functionality subject to preprocessing raw data into staging tables.

Key Analytics Components of YODA

Demand Forecast

YODA forecasts the number of customers who will buy a product, at what time, and at what price. With this information, the optimization model can estimate demand for each product and adjust prices to maximize the revenue generated across those products. YODA generates a booking curve per week for public fare types for all available products. YODA then aggregates demand to ensure it is at the level of granularity that the optimization model requires (i.e., by submeta, source market, time interval, and product type). The demand that YODA addresses constitutes most of the individual bookings that Carnival takes for its cruises and represents the audience that will want this product at a given price.

YODA forecasts assume passengers are in lower berths. An upper berth might be an upper bunk in a cabin, a baby crib, or a foldaway double bed. A typical cabin accommodates two adults; however, some cabins accommodate up to four guests (e.g., a family traveling with children). A ship can also include a small number of dedicated single cabins and two-bed cabins with single occupancy. Because guests must always book all lower berths in a cabin, YODA relates the bookings to the lower-berth capacity.

YODA's demand forecast process identifies the relevant history, which forms the basis of the initial forecast. Further adjustments to the forecast take into account differences in the attributes of historical and current voyages, as well as differences in year-over-year market performance.

Identification of Relevant History

Given the complexity of the cruise product (e.g., multiple itineraries, durations, trades, ships, cabins, and fares), YODA examines historical booking data to identify history that is relevant to forecasting demand.

YODA uses classification and regression trees (CART) at a brand and trade level, as well as business rules to identify the voyage attributes (e.g., subtrade, duration, time of year, and ship class) that best predict the booking-curve behavior on cruises in that trade, based on the historical booking data. CART is a simple but powerful machine learning method used in constructing predictions from data (Krzywinski and Altman 2017). The CART models recursively partition voyage history into subsets based on similarity of behavior and provide relative weights across different voyage attributes. The output is used to find a weighted set of like cruises to use as history for forecasting future cruises.

YODA removes outlier historical voyages by calculating median booking-curve confidence intervals and identifying voyages whose booking curves fall outside the confidence intervals for more than a specified amount of time. These historical sailing matches are then reviewed in conjunction with the RM teams who are ultimately responsible for reviewing and accepting price changes. In some cases, the RM teams override the selected historical voyages, for example, if the number of passengers and/or revenue had been affected by a significant event like an onboard celebrity.

INFORMS Journal on Applied Analytics, 2021, vol. 51, no. 1, pp. 26–41, © 2021 INFORMS

This combination of impartial data science rigor and key insights from the business allow the demand forecast to reach a level of accuracy unattainable with either method alone.

Unconstraining Booking Curves and Making Adjustments Based on Future Voyage Attributes

Once YODA identifies valid history for a future voyage, it makes an adjustment to account for the differences between historical and future voyages. YODA first unconstrains the historical booking curves by estimating latent demand, as we describe in the phase 1 overview.

Because of changes to yearly deployment of cruise ships by Carnival and other cruise lines (OCLs), demand for a product within trades will shift in relation to industry capacity. If any cruise line changes its itineraries or releases a new ship into a trade, then the capacity of available lower berth-days (ALBDs) will change, impacting demand for Carnival's products.

YODA calculates average Carnival and OCL ALBDs and uses regression techniques to determine capacity normalization factors between ALBDs and prices paid for different cabin types and products at various times of the year. These factors are then applied against historical booking curves for like cruises to generate a set of unconstrained historical booking curves at capacity-agnostic prices. These unconstrained and capacity-adjusted historical curves are combined into a weighted-average booking curve at a reference price (i.e., historical average price) and smoothed using exponential smoothing to control for variabilities created by historical behavior.

Since 2018, these capacity normalization factors generated in YODA have also been used continuously as key components in determining the optimal ship deployments within the Carnival deployment optimizer. The Carnival deployment optimizer is a system independent of YODA, which was created to assist with the yearly deployment of Carnival cruises.

Adjusting Voyage Forecasts Based on Performance and Market Dynamics

A normalization process tunes the demand forecast by reconciling current voyage and market performance with historical performance. A year-over-year market module determines the amount of Carnival's additional (or reduced) demand in real terms—each historical cruise will have its own demand curve with bookings taken at different prices, often where different pricing and promotion strategies have been used from year to year. The module uses price elasticities of demand to adjust for the differences in demand due to variations in pricing and applies the same price normalization effect to the bookings we have taken so far on a future cruise.

YODA then incorporates recent trend models into the forecast. The recent-trend model dynamically determines the number of past weeks to use in comparing forecasts with actuals and computes a difference factor at the appropriate level, which it then applies to a number of future weeks (again dynamically determined). Because cruises can be on sale two years before departure, the like-cruises model can select, for example, a cruise departing one year from now as a like cruise to represent a cruise departing two years from departure. In this case, we use the history we have gathered so far from next year's cruise to update our forecast for the cruise that departs the following year. To ensure that the history remains as up-to-date as possible, YODA then makes a second pass to reapply the trend adjustment from the earlier future cruise to the later future cruise.

Booking Retention

Some passengers who book a cruise will not travel on that cruise. The reasons vary from financial limitations, unavoidable changes in plans (e.g., because of family health reasons), or just changing one's mind. The conditions and age of a booking will dictate the penalties that apply for a cancellation. YODA uses CART models to determine the attributes that are important in predicting retention. A booking-level survival forecast from the booking date to the cruise date updates that probability as any detail about the booking changes (e.g., deposit payments, upgrades, additional purchases). These retention probabilities are computed for each booking to enable Carnival to identify high-value bookings and guests with a higher risk of cancellation so that it can take actions to increase the probability of retaining those bookings.

YODA needs to determine the number of bookings to take during the remaining booking window to ensure that when the cruises sail, they are as full as possible, because more highly occupied ships drive additional onboard revenue. As a result, YODA optimizes using a *net-net* forecast; that is, it applies the retention probability to the bookings already on the books at that point in time. It also applies the forecasted retention probability to the bookings yet to be made.

Time Intervals

One of the fundamental requirements of YODA is to produce a recommended price strategy, generating an optimal price both now and at strategic points in the future. These points, such as the final-payment due date (for passengers who paid an initial deposit on their bookings, this is the date on which the balance is due), are used to split the booking horizon into time intervals. The intervals are derived by dividing the historical booking curve before the final-payment due date into equal proportions of demand using the number of time intervals defined by each brand.

These time intervals are applied to each voyage, base market, fare type, and meta so that each bookable product has multiple price points. Demand is aggregated for optimization into these time intervals so that YODA can set a recommended price for each time period and hence give an optimal price profile across the life of the cruise.

Market Response Modeling

YODA uses an exponential market response model (MRM) to estimate price elasticities. This allows the optimization model to use how demand reacts to changes in price and then adjust prices to set the optimal price and demand for each product across the time intervals. By taking multiple fare types into consideration, YODA ensures that more passengers are given a price that best meets their needs and elasticity to price.

YODA's fare types integrate with pricing structures and relationships (including benchmarks) that facilitate the RM system's interaction with price code structures (promos) within POLAR. Each booking is associated with a fare type depending on the promos attached to it. Within each fare type, there is one benchmark fare on which other fares depend, based on a set relationship (e.g., a price for loyalty guests may always be 10% less than the benchmark fare). All historical prices paid are translated back to their benchmarks. The benchmark fares are used in the MRM model to avoid creating bias on the price-change effect by different promos.

Many RM applications use a linear market response model. However, phase 1 exploratory work showed that an exponential curve best represented cruise customers' behavior. An exponential curve captures the extreme cases (e.g., extremely high or low prices) much better than a linear curve does. YODA uses CART models to cluster together similar bookings in terms of their response to price and fits an exponential curve (Figure 4).

Demand for Groups of Guests

Some Carnival brands sell significant portions of cruises (up to 100%) at fixed rates that they prenegotiate with travel agents, with limits set on the number of bookings sold at a given rate. Travel agents are likely to be optimistic about the number of bookings they think they can sell at that rate, and they often request a contract for a set number of cabins with the knowledge that they may not sell all the cabins.

The groups arranged by travel agents are generally not as responsive to prices as are passengers for regular bookings; consequently, Carnival treats this demand stream differently. Group bookings are excluded from historical demand, and for groups on current voyages, YODA calculates their net-net forecasts

Figure 4. (Color online) MRM Curve Shows the Exponential Curve Fit of Capacity-Normalized Demand Against Price per Day for a Given Market Segment

based on historical group behavior. Group bookings, however, can be cancelled in the same way that regular bookings can. For most groups, YODA uses a CART model to estimate the conversion rate from that group based on the group's contracted size. Conversion rate captures the proportion of the group size, sold to agents, that is booked and results in a booking that specifies the cruise being taken. YODA includes factors such as group type, market, and travel agency details, as well as the size of the group allocation request and cruise duration.

Optimization

We successfully implemented a joint pricing and inventory optimization process, a novel approach in RM, which can handle multiple cabin occupancies and multiple capacity-constraint definitions—a problem unique to the cruise industry. The objective is to maximize net ticket and ancillary passenger revenue minus costs subject to inventory availability and pricing constraints, while incorporating the risk of guests cancelling and rebooking at a lower rate if the price decreases, as well as any additional contribution from onboard revenue (Figure 5). The objective function is a quadratic concave function representing the contribution of the network of voyages in the form of revenue minus cost in the currency in which that brand operates (e.g., British pound sterling for P&O UK or U.S. dollars for Princess Cruises). Factors included are commission cost (i.e., a percentage of revenue that is paid to a travel agent on securing a booking), onboard cost, air cost, cruise tour cost (e.g., Denali National Park tours in Alaska) for all cruises, fare types, submetas, berth types, products, base markets, and time intervals. YODA achieves this using a quadratic program (QP), which leverages price elasticity of demand, demand forecast (based on reference prices), and projected availability as primary inputs. The objective is to simultaneously optimize pricing plans and inventory allocations (including upgrades) throughout the booking window across all overlapping bookable cruises on a ship and meta by setting duration controls at a meta/sailing level. YODA's QP model also factors in the risk from *retro-dilution* or *retro*, which is the risk of guests cancelling their earlier booking and rebooking the same voyage at a lower fare. We only partially describe the formulation of the optimization model in the appendix because the full formulation is proprietary.

Figure 5. (Color online) YODA's Joint Price and Inventory Optimization Model Takes Multiple Inputs and Makes Trade-offs Across Multiple Voyages

YODA's QP optimization engine runs independently for all six original brands, plus any new brands joining YODA (e.g., AIDA), regardless of their network configurations, time intervals, products, berths, fares, and base-market configurations.

A ship can have various embarkation ports leading to a variety of possible cruises; modeling network configurations allows the optimal allocation of inventory across these cruises. We list key aspects of YODA here.

- An innovative feature of YODA is that it divides the booking horizon into multiple time intervals. The model then simultaneously recommends inventory allocations and prices across these time intervals rather than recommending only the current price. It also incorporates constraints linking the pricing across these time intervals if the business wants to enforce constraints such as increasing prices only as the voyage's sailing date gets closer.
- Inventory allocations, based upon the recommendations of the optimization model, are adjusted throughout the booking horizon. As customer demand materializes and the demand forecast changes, the inventory allocation recommendations also change. The analysts have the ability to freeze inventory allocations for a specified period of time during the booking horizon such that only price is being optimized.
- Options are available for guests to buy cruise tours or book air flights with the cruise; the optimization model recommends a bundled price across these products and can simultaneously generate price recommendations for both unbundled and bundled products.
- The optimization model recommends pricing across the multiple fare types and upper/lower berths, and forces minimum and maximum price differentials between them.
- YODA forecasts demand and estimates price elasticities across different base markets, allowing the optimization model to recommend prices that can differ across these base markets after considering exchange rates. This strategy of optimizing price based upon the characteristics of the market segments

was introduced by Agrawal and Ferguson (2007) in the context of bid prices, which we have extended to our application.

• YODA can use the outputs of the exponential elasticity model by finding the tangent line at the reference price as the linear approximation to the fitted exponential curve.

Through network optimization, YODA recommends the number of cabins to be allocated to each bookable voyage within the network. Additionally, YODA has nested capacity constraints that allow for overbooking across the categories of cabins.

YODA's QP incorporates advanced RM strategies beyond elasticity-driven price recommendations. For example, YODA avoids the booking of high-occupancy (i.e., four-berth) cabins by a lesser number (i.e., one or two) of guests, thereby enabling upgrades to more premium cabins. To ensure the price integrity and business reasonability of price recommendations and to address specific brand nuances, the price constraints, additional revenues, some fixed and variable costs, and associated business rules are incorporated into the optimization model using a complex setup that accounts for cabin type, cruise tour bundling, and air flight bundling. YODA can also optimize, based on optional minimum price constraints, to represent the tradeoffs between selling cabins publicly and holding space for groups, as negotiated with travel agents.

YODA also must ensure that our solution satisfies both our lower-berth capacity constraints (i.e., the number of physical cabins on the ship) and the passenger safety-limit constraint (i.e., number of guests permitted onboard) while maximizing revenue. YODA accomplishes this by using an average occupancy rate, which allows it to optimize for both upper and lower berths.

Notable constraints of the optimization model formulation are as follows.

• The model allows *elastic* inventory allocations across submetas of a cruise (e.g., overbooking) or across cruises of a network while considering the retention effects of existing bookings and a ship's total passenger safety limit. Similar to the artificial intelligence (AI) Tetris game, YODA selects and allocates blocks of inventory that best accommodate its forecasted demand patterns within a voyage and across overlapping voyages, thus enabling dynamic inventory reallocation of ship capacity.

• It integrates complex pricing architecture requirements to support the implementation of brand-specific pricing strategies across time intervals, metas, fare types, product types, and cruises.

Impacts of Solutions

Revenue Uplift

At the conclusion of the testing period that we describe in *Phase 2: Build and Market Test Prototype Solution*, estimates of annual revenue uplift ranged from (+1.5% to +2.5%), with results varying based on the comparison methodology used and by brand. Because the test period did not cover an entire booking curve, we needed to make an assumption on how the remaining demand would materialize. To make a fair comparison between the test and control voyages, YODA used the optimal prices and the expected demand at optimal prices for both the test and control voyages. This means that the revenue uplift range we show above is a conservative estimate because the testing approach benefited the control voyages as much as the test voyages. Because of the positive results of the test, more voyages were continually moved to YOGA management, thus preventing an ongoing measurement of uplift.

YODA has generated price and inventory allocation recommendations for all on-sale cruises of the six brands since December 2017, and each brand uses its recommendations. The efficiency of the price update process has resulted in saving thousands of people-hours since that time. In addition, based on a controlled trial, we estimated that retention increased by approximately 1.5% as a result of improved customer communications.

Change Management

YODA ushered in a new mindset about revenue and inventory management at Carnival. A user guide, the playbook, is comprehensive in that it describes all aspects of the system from how often prices should be reviewed to the screens the analyst accesses to complete a review.

The playbook emphasizes that YODA is designed to recommend a price change in response to changes in demand or some external factors. When YODA recommends a price increase, analysts can quickly take that action. However, when YODA recommends a price decrease, analysts should also consider other actions. Instead of driving demand through price, they can work with other commercial teams to offer sales incentives, increase marketing, or create promotions. This shift in thinking resulted in a major change in managing voyages.

Challenging Legacy Processes

In addition to the revenue benefits highlighted previously, YODA has driven renewed strategic conversations within the RM teams in each of the six brands by giving the data science teams and analysts the ability to challenge long-standing principles of cruise RM using YODA outputs and associated models. For example, they rigorously challenged the high-to-low nature of pricing in some markets and thus enabled further conversations on marketing and price positioning.

Recognition and Encouraging Collaboration

YODA has set a new standard for cross-brand projects in Carnival Corporation. Not only did the project complete on time and below budget, but it also fostered a level of collaboration never seen previously among the brands. Teams now regularly reach out to each other to discuss how to create promotions within the reservations system, how to efficiently load pricing and planned changes to the pricing strategy, and how to share learnings and successes on how particular promotions and strategic decisions have impacted a particular market. In addition, many of the early champions of the YODA system and project within the brands have been promoted into higher roles within the business.

Fostering Analyst Development and Job Satisfaction

Greater responsibility and authority are owned at the analyst level now, allowing the analysts to have greater control of their day-to-day work. This not only heightens analysts' abilities to identify and interpret insights, but it also improves their job satisfaction and provides better career progression. Tactical decision making has improved because analysts have the time to look at details and make crucial decisions. Previously, senior-level executives made many pricing decisions based on their intuition and experience. Company confidence in the combination of YODA and the analysts means that many senior decision makers no longer feel that they must attend pricing meetings, thus allowing them to focus on more strategic issues.

Improving and Simplifying Product Offerings for Guests

Internal conversations about the value and use of value-adds, such as drink packages and free onboard spending money (i.e., onboard credit), have matured. For some of our brands, decisions about price discounts previously fell outside of the RM processes. YODA benchmarks now control these discounts, and other fares are priced relative to this benchmark fare; this in turn provides greater pricing consistency to our guests. Additional payment and package options are now available to consumers, allowing them to better customize their holidays and pay only for the services they value.

Using YODA Science Across the Business

Both P&O Australia and Princess Cruises have used retention models from YODA. In some cases, the brands have modified the models slightly to include additional variables. For example, to satisfy the European Union's General Data Protection Regulation (GDPR), YODA contains no personally identifiable information. These models are used to identify booking-level retention probabilities at a very granular level and to identify bookings that have a low probability of retention; brands can proactively take steps to encourage these guests to honor their bookings. The brands may contact guests by various means, depending on their risk of cancellation. As we highlighted previously, in addition to the 1.5%–2.5% revenue uplift, A/B tests of bookings of contacted versus noncontacted guests demonstrated an increased retention of approximately 1.5% in a controlled trial.

Growth and Integration of the Operations Research Teams

At the beginning of the project, only four data scientists were employed by the six brands in their operations research teams. These teams now include at least 15 data scientists plus other associated roles, with involvement in areas such as RM, marketing, onboard revenue, website optimization, guest sourcing, and deployment.

Single Source of Truth for Analytics

YODA represents an essential consolidated database for ongoing and in-depth data analytics and a common data set and framework for RM reporting. Its insights are used to support onboard revenue decisions. Daily snapshots provide a depth of history not available previously, common definitions and platforms allow the transfer of analysis across brands, and reports to key industry groups are now based on the YODA database. This facilitates rapid cross-brand information sharing and fosters a culture of iterative development and peer review.

Cross-Industry Impact

In addition to being easily portable to other cruise lines, as highlighted in *Phase 3*: *Enterprise Solution, Design, Beta Testing, and Rollout*, models developed in YODA are valuable and applicable for other industries. The joint price and inventory optimization model is portable across the travel and hospitality industries. The demand forecast learning algorithms we developed are applicable across any industry that uses RM, specifically casinos and resorts where demand is special-event driven and does not follow a typical time-series pattern. The exponential market response model is relevant in industries where the relationship between price and demand is nonlinear.

Acknowledgments

The role played by the executive committees was critical in ensuring the project remained strategically well aligned with business requirements. Key members of these senior executive committees were as follows: Josh Leibowitz (All Brands Group), Deanna Austin (Princess Cruises), Beth Bodensteiner (Holland America Line), Matt Rutherford (P&O Australia), Nigel Esdale (Cunard and P&O UK), and Tobias Waack (AIDA). The following people were instrumental in the development of YODA and we publicly acknowledge their contributions: Corrado Sala, Kristina Kaylen, Andrea Torassa, Stefano Riva, Laura Whittaker-Rawlinson, Mitchell Peacock, and Mike Planicka (All Brands Group); Mike Miele, Ravi Regalla, Aashish Seethamraju, Jathin Chandran, Nancy Babu, Puspanjali (Anjali) Kumari, and Tsipi Yaffe (Shared Services); Caroline Handisides, Justin Beck, Jenny Morgan, Melinda Urban, and Malgorzata Pietka (P&O Cruises Australia); Richard Warne, Mike Whales, Chris Clark, Stephanie Cross, Michael Winterbottom, Prakhar Singh, Jeannine Mellor, John Harvey, Peter Vermeulen, Hattie Hall, Clement Li, Marco Tufail, Alex Anderson, and Stefano Riva (Cunard and P&O UK); Dean Apy, Lisa Syme, Fina Lam, Yanqi Xu, Justin Beck, Wei Xie, Andrew Poduska, Bernard Niu, Peng Yang, Cole Sanders, and Danny Meecham (Princess Cruises); Sean Brennan, Paul Grigsby, Greg Vogel, John Harvey, Brian McDowell, Fred Gardner, Meggan French, Alejandro Avalos Mar, Chuanfeng (Ashley) Yang, Alexey Lagerberg, and Elisa Ahern (Seabourn and Holland America Line); Daniela Versick, Maik Bormann, Christian Uterwedde, Norman Wilken, and Hannes Schmeiduch (AIDA); Matt Busch, Jon Higbie, Steven Moy, Stan Ward, Kristina Kaylen, Dan Iliescu, Pratik Mital, Jordan Avery, Kais Hadj-Taieb, Drew Yao, Chuan Mao, Kristen Letchworth, Andrew Tippey, Justin Ji, Elaine Zhao, Bobby Falconer, Ben Johnson, and Peter Chao (Revenue Analytics).

Appendix. Formulation of Optimization Model

This appendix describes some of the essential elements of a quadratic programming model, which recommends prices for cruises within a brand's network of sailings (i.e., cruises) to maximize net ticket and ancillary net passenger revenue. Furthermore, when a ship may be used by multiple cruises on a same day, the model determines the optimal allocation of the ship's cabin inventory to cruises. Definitions are presented in Table A.1.

We can then formulate the model as follows:

$$\text{Maximize} \sum_{s,f,c,b,k,m,t} \left(ucd_{sfcbkmt} - a_{sfcbkmt}\right) \cdot x_{sfcbkmt} \tag{A.1}$$

subject to

$$ucd_{sfcbkmt} = df_{sfcbkmt}\left(1 + e_{sfcbkmt}\left(\frac{x_{sfcbkmt}}{p_{sfcbkmt}} - 1\right)\right), \tag{A.2}$$

$$\sum_{c' \le c,\, s \in S(l), f, m, t} \left(ucd_{sfc'bkmt} - a_{sfcbkmt}\right) \le \sum_{c' \le c} rc_{lc'b}, \tag{A.3}$$

$$ucd_{sfcbkmt} \ge 0,\ x_{sfcbkmt} \ge 0,\ a_{sfcbkmy} \ge 0.$$

Objective function (A.1) maximizes net revenue as the product of net price x and constrained expected demand, which is defined as the difference between unconstrained expected demand ucd and excess demand a. Constraint (A.2) uses the elasticity equation to link net price and unconstrained demand. Constraint (A.3) ensures nested constrained demand does not exceed capacity for each combination (i.e., leg, submeta, berth). For each submeta, the model explicitly allocates the berth inventory to the voyages that use them, and these allocations are among the model outputs reviewed (and updated) by the brand's pricing analysts.

Table A.1. Definitions

Panel A: Sets and indices	
L	All sailing legs of a network: $l \in L$
S	All sailings of a network: $s \in S$
$S(l)$	Subset of sailings that includes leg l: $l \in L$
F	Set of fare types, $\{1,2,\cdots,n_F(s)\}$, ordered with 1 as the highest fare type: $f \in F$
C	Set of submetas, $\{1,2,\cdots,n_C(s)\}$, with 1 as the highest category group: $c \in C$
B	Set of berth types (i.e., $\{Lower, Upper\}$): $b \in B$
M	Base market: $m \in M$
$T(s)$	Time intervals of booking window for sailing $t \in T(s)$
K	Set of offered inventory products, $\{CruiseOnly\}, \{Cruise + Cruise\ Tour\}, \{Cruise + Air\ Travel\}$: $k \in K$
$S(k)$	Set of sailings that use product k: $s \in S(k)$
Panel B: Key terms	
Price elasticity	Change in demand expected for a given change in price, taken as the slope of the tangent line to the MRM curve at the reference price
Remaining demand	Expected demand at the current price (reference price for future time intervals)
Products	The set of different bookable options for a cruise; for example, a cruise can be booked including/excluding air travel and including/excluding land tours (Alaska only); each is priced individually
Sailing leg	Collection of voyages into a *ship day*; captures all bookable voyages that, if booked, would be available on a given ship on a particular sailing day
Capacity	Physical lower-berth capacity of the ship
Available capacity	Lower-berth capacity of the ship still available for booking, after factoring in retention of existing booking

Table A.1. Definitions. (Continued)

Optimizable capacity	Lower of (1) the ship's available capacity, or (2) the remaining number of passengers the ship is allowed to accommodate before reaching its passenger safety limit (often determined by lifeboat capacity)
Excess demand	Amount by which demand exceeds optimizable capacity
Unconstrained expected demand	Expected demand at the optimal price recommended by YODA; this is unconstrained to allow for price rules to be satisfied and for demand to be in excess of capacity at the submeta level, to allow for upgrades, and for reallocation of capacity allocated to each bookable voyage
Nested constrained demand	Unconstrained expected demand constrained by capacity, summed over all submetas up to and including the current submeta, ordered by the submeta hierarchy; this permits the optimization to recommend overbooking and upgrades upward only on the ship.
	Panel C: Input parameters
$p_{sfcbkmt}$	Net current price (ticketed fare less value adds and onboard credits) for sailing s, fare type f, submeta c, berth b, product k, base market m, and time interval t (for future time intervals, we use the reference price instead of current price)
$e_{sfcbkmt}$	Price elasticity at price $p_{sfcbkmt}$ for sailing s, fare type f, submeta c, berth b, product k, base market m, and time interval t
rc_{lcb}	Remaining available capacity of sailing leg ?, submeta c, and berth b
$df_{sfcbkmt}$	Remaining demand forecast (not yet materialized) for sailing s, fare type f, submeta c, berth b, product k, base market m, and time interval t
	Panel D: Decision variables
$x_{sfcbkmt}$	Optimal net fare (ticketed fare less value adds and onboard credits) for sailing s, fare type f, submeta c, berth b, product k, base market m, and time interval t (nonnegative)
$a_{sfcbkmt}$	Excess demand above optimizable capacity for sailing s, fare type f, submeta c, berth b, product k, base market m, and time interval t (nonnegative)
$ucd_{sfcbkmt}$	Expected demand at optimal net fare for sailing s, fare type f, submeta c, berth b, product k, base market m, and time interval t (nonnegative and can exceed remaining available capacity; hence, it is unconstrained)

The model includes many other objective function components, decision variables, and constraints to ensure business-reasonable pricing across time intervals, fare types, submetas, berth types, and product types (as we describe in the main body of the paper). Examples of these are as follows:

- Risk associated with guests cancelling and rebooking at a lower rate if the price decreases;
- Additional contribution from on-board revenue;
- Estimated commission costs (i.e., a percentage of revenue that is paid to a travel agent on securing a booking);
- Air costs (and associated capacity constraints by airport) for cruises in which airline seats are chartered to transport the guests to the ship; and
- Cruise tour costs (and associated individual tour capacity constraints) for all Alaska cruises that offer the option of a combined cruise and tour of Denali National Park.

Not all components of the model are depicted in the previous formulation because the full optimization is considered proprietary. This problem is formulated in such a way that CPLEX can always find a unique solution.

References

Agrawal V, Ferguson M (2007) Bid-response models for customised pricing. *J. Revenue Pricing Managemeent* 6(3):212–228.

Ayvaz-Cavdaroglu N, Gauri DK, Webster S (2019) Empirical evidence of revenue management in the cruise line industry. *J. Travel Res.* 58(1):104–120.

Belobaba PP (1987) Air travel demand and airline seat inventory management. PhD dissertation, Massachusetts Institute of Technology, Cambridge.

Cross RG (1997) *Revenue Management: Hard-Core Tactics for Market Domination* (Broadway Books, New York).

Cross RG, Higbie JA, Cross ZN (2010) Milestones in the application of analytical pricing and revenue management. *J. Revenue Pricing Management* 10(1):8–18.

Cruise Lines International Association (2019) 2020 State of the cruise industry outlook. Accessed June 24, 2020, https://cruising.org/-/media/research-updates/research/state-of-the-cruise-industry.pdf.

Fiig T, Isler K, Hopperstad C, Belobaba P (2009) Optimization of mixed fare structures: Theory and applications. *J. Revenue Pricing Management* 9(1–2):152–170.

Guillen J, Ruiz P, Dellepiane U, Maccarrone L, Maccioni R, Pinzuti A, Procacci E (2019) Europcar integrates forecasting, simulation, and optimization techniques in a capacity and revenue management system. *INFORMS J. Appl. Analytics* 49(1):40–51.

Hormby S, Morrison J, Dave P, Meyers M, Tenca T (2010) Marriott International increases revenue by implementing a group pricing optimizer. *Interfaces* 40(1):47–57.

Krzywinski M, Altman N (2017) Classification and regression trees. *Nat. Methods* 14:757–758.

Kuyumcu A, Yildirim U, Hyde A, Shanaberger S, Hsiao K, Donahoe S, Wu S, Murray M, Maron MB (2018) Revenue management delivers significant revenue lift for holiday retirement. *Interfaces* 48(1):7–23.

Smith BC, Leimkuhler JF, Darrow RM (1992) Yield management at American Airlines. *Interfaces* 22(1):8–31.

Talluri KT, van Ryzin GJ (2004) *The Theory and Practice of Revenue Management* (Springer, New York).

Justin Beck graduated with a BSc in mathematics and physics. He later graduated from the University of South Australia with a PhD in mathematics, specializing in game theory. He spent a number of years working in military operations for the Australian Department of Defence, and also spent time at RAND Corporation. He joined Carnival Australia in 2015 and was the lead operations research team

member for P&O Australia before promotion to the business leadership team.

John Harvey has worked for Holland America Line as director of revenue science since May 2019, and his role closely involves conception, validation, and introduction of new levels of scientific and mathematical method into forecasting and optimization processes, as well as driving new operations research opportunities within cruise. He received a master's degree in operations research from the University of Warwick, United Kingdom and, since 2008, has worked across the travel industry.

Kristina Kaylen is the director of Group Revenue Performance on the All Brands Group Strategy Team within Carnival Corporation. Before Carnival, she worked in revenue management consulting, leading operations research and strategy teams to build transformational systems in cruise, hospitality, automotive, and retail. Kristina received a BS in industrial engineering at Georgia Institute of Technology.

Corrado Sala is senior vice president of Group Strategy at Carnival Corporation. In the YODA program, he built, led, and inspired a diverse talented team to work together, translating latest business priorities into scientific optimization processes, promoting an overall governance based on change management. In his fast-track career at McKinsey & Company between 2005 and 2014, he helped Fortune 500 companies and private-equity owned companies in step changing commercial performance.

Melinda Urban is a senior data engineer at Carnival Australia. During and after completion of her degree in economics in Hungary, she worked in the telecommunications industry for a major European provider. In the initial development phase, her responsibilities involved data validation and testing of YODA's science modules. She is the lead member of the Science team representing P&O Australia, ensuring YODA keeps performing at its best.

Peter Vermeulen is currently a senior data scientist and the YODA lead scientist for P&O Cruises and Cunard. He received an MSc in operational research with finance at the University of Southampton. Peter joined the YODA project from an operations research role at the Automobile Association, driving significant development, integration, and adoption within the project. He leads a team of data scientists responsible for adoption, performance, and development of YODA.

Norman Wilken is the lead for the Yield Science and Projects team within AIDA Cruises. He has worked within the company for 2 years and is leading a team responsible for the development and implementation of new technical solutions and algorithms to help increase ticket revenue within the brand. Norman developed his scientific skillset during a PhD program at Chalmers Technical University and expanded within work in the automotive industry.

Wei Xie is a senior manager of revenue data science at Princess Cruises. His team develops the revenue management system (RMS), supports pricing teams to use RMS, and supplies data science expertise to various internal business teams. He has 10+ years in revenue management and data science and industry experiences in cruise lines, airlines, and finance. He received his PhD degree from University of Illinois at Urbana–Champaign and BS degree from Fudan University in China.

Dan Iliescu is a senior director of science at Revenue Analytics. His work with the firm includes development, configuration, and testing of next-generation revenue management systems for travel and hospitality industries. His analytical expertise covers artificial intelligence–powered algorithms for B2B sales and B2C availability and pricing controls. He has 10+ years of experience with multiple travel and hospitality clients creating innovative revenue management cloud applications.

Pratik Mital leads the science team in Rail and Cruise at Revenue Analytics. He has 8 years of experience in applying machine learning and optimization techniques across various fields such as cruise, hospitality, rail, and airlines. He has conducted cutting-edge research in the field of pricing and revenue management, as well as strategic design of supply chains. He has also applied operations research techniques to model reforms in the K-12 education system in the United States.