Research Article

Dynamic pricing – The next revolution in RM?

Received (in revised form): 7th February 2016

Thomas Fiig[a], Oriana Goyons[b], Robin Adelving[b] and Barry Smith[c]

[a]Amadeus Airline IT, Kastrup, Denmark; [b]Amadeus Airline IT, Amadeus S.A.S., Sophia Antipolis Cedex, France; and [c]Barry C Smith LLC, Dallas, USA

Thomas Fiig is director, chief scientist at Amadeus, where he is responsible for revenue management strategy and scientific methodologies. He holds a PhD in Theoretical Physics and Mathematics and a BA in Finance from the University of Copenhagen, Denmark. He has published several articles, recently focused on methodologies for origin-and-destination forecasting and optimization of simplified fare structures. He serves on the editorial board of the *Journal of Revenue and Pricing Management.*

Oriana Goyons is a product definition analyst in the revenue management area at Amadeus. Currently, she is constructing the dynamic pricing product. She earned her master's degree in Operations Research from INPG Institut National Polytechnique and her engineering degree from ENSIMAG in Grenoble, France.

Robin Adelving is a software development engineer at Amadeus, and a main contributor to the development of the Altéa Benchmarking Engine (ABE) simulation tool. Currently, he works in revenue management and operations research. He holds an engineering degree in Mathematics and Computer Science from ENSIMAG in Grenoble, France.

Barry Smith is a freelance engineer specializing in pricing, revenue management and air transportation scheduling. While at American Airlines and Sabre, he developed many of the revenue management techniques used throughout the airline industry and pioneered the application of revenue management techniques in other industries. He holds patents in reservation system design, scheduling, pricing, operations and online retailing. He is an AGIFORS past president and Fellow. He led a team that received the Edelman Award and INFORMS Revenue Management Section Prize for their work on revenue management at American Airlines.

Correspondence: Thomas Fiig, Amadeus Airline IT, Lufthavnsboulevarden 14, Kastrup DK-2770, Denmark.
E-mail: Thomas.Fiig@Sca.Amadeus.com

ABSTRACT Existing revenue management systems (RMS) base their recommendations on historic observations and do not explicitly consider competition. This means that RMS recommendations often are not appropriate for real-time competitive situations. *Dynamic pricing* (DP) is an extension of RMS that dynamically calculates the optimal price, taking into account the airline's strategy, customer-specific information and real-time alternative offerings. By optimizing the contribution within the shopping session, DP has a more current and detailed view of demand and can improve RMS performance. We investigated the performance of DP using two simulators, Altéa Benchmarking Engine and Passenger Origin Destination Simulator and demonstrate that DP can deliver substantial revenue benefits with no modification to existing revenue management (RM) processes. However, the deployment of DP into the airline distribution process will be a challenge, because it affects all shopping and downstream processes, such as ticketing, servicing, revenue accounting, RM and interline settlement, that rely on information from existing fare aggregators. Nevertheless, the potential benefits of DP are so compelling that we believe the effort to bring this technology into practice is warranted.
Journal of Revenue and Pricing Management advance online publication, 3 June 2016; doi:10.1057/rpm.2016.28

Keywords: dynamic pricing (DP); distribution; revenue management system (RMS); global distribution systems (GDS); Passenger Origin Destination Simulator (PODS); Altéa Benchmarking Engine (ABE)

INTRODUCTION AND MOTIVATION

The Internet has fundamentally transformed how airline customers shop, bringing unprecedented price transparency to the marketplace and making revenue management (RM) decisions more dependent on competition. Existing revenue management systems (RMS), however, are blind to real-time competitive situations. The RMS uses historical data to perform forecasting and optimization and then relies on an offline upload of bid prices to the online Inventory system. RMS takes competition into account only indirectly (and in an averaged way) through its effect on bid prices. Hence, an RMS will provide identical recommendations regardless of the real-time competitive scenario.

The current RMS paradigm, which relies on the process of fare filing and then having RMS opening or closing fares, is limited in two ways. First, fares are static and published before being offered publicly. Second, these fares and availability controls are applied in an identical fashion across multiple shopping sessions, during which the optimal price may vary significantly.

We believe that in the current business environment, *dynamic pricing* (DP) can improve RMS performance. DP provides benefits in the following areas:

Real-time responsiveness: The choice set within a shopping session depends on competitors' schedules and available prices. While schedules and prices are relatively static, availability changes frequently. DP allows airlines to adjust prices in real time during the shopping session (without the need to file fares in advance) based on the competitors' offers.

Customer segmentation: DP can observe customer characteristics during the shopping process and may improve its estimate of a customer's willingness to pay (WTP), expectations of service quality or preferred time of day. With this ability, the airline can improve its product offerings to the customer.

Pricing precision: Even though airlines may publish prices several times per day, they offer a relatively fixed set of prices for each market, with gaps between prices. The best prices for an airline's products during a given shopping session may lie between or outside of the range of published fares. DP improves pricing precision by offering a continuous price.

Industry-changing

DP will revolutionize RM in two significant ways. First, DP will change RMS from a purely offline process involving processing of historical data into a dynamic, real-time pricing engine. Second, DP will disrupt the current distribution processes. With DP, fares and taxes filed to fare aggregators such as ATPCO are no longer sufficient for price computation. This affects all shopping and downstream processes (for example, ticketing, servicing, revenue accounting, RM and interline settlement) that rely on existing information from fare aggregators.

While DP seemed farfetched in the past, the advent of the New Distribution Capability standard has made DP both realistic and desirable for airlines.

Contribution

To our knowledge, no published work provides an investigation of the practical aspects, deployment or operation of DP, nor are we aware of any DP implementation within the airline industry.

Our goal is to investigate DP's behavior and anticipate its likely revenue performance to supplement existing theoretical work. We demonstrate consistent benefits for early adopters of DP through simple examples and realistic simulations. We also show that DP can benefit multiple airlines in a single market. Significant work remains to be done to bring DP into everyday use. In this article we provide context and motivation for these next steps.

Organization

This article is organized as follows: After a literature review, we define DP and discuss the

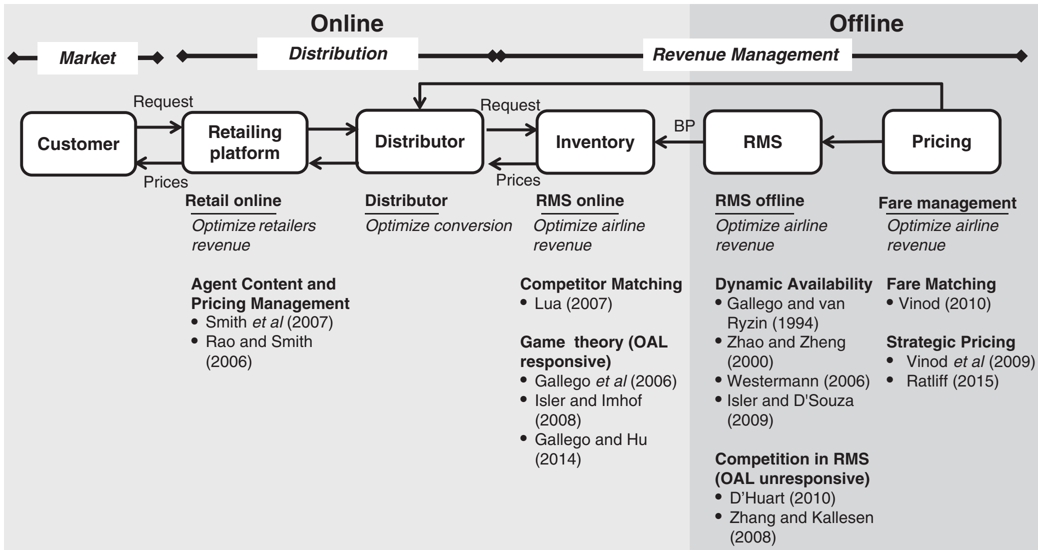

Figure 1: Literature review – Competitive awareness.

fundamental difference between DP and RMS. This is followed by the presentation of revenue results from two simulators, Altéa Benchmarking Engine (ABE) and Passenger Origin Destination Simulator (PODS). Finally, we discuss limitations and ideas for future research, and present our conclusions.

LITERATURE REVIEW

Figure 1 shows how a customer request for an itinerary is passed through the retailer, distributor and to the airline's Inventory system for evaluation. The Inventory system performs the online RMS control that calculates availability, which, via the pricing system, determines price quotes that are communicated back to the customer. RMS offline processes, which include forecasting and optimization, calculate bid prices that are used in online RMS control. Pricing is used both by RMS for yield estimates and by distributors (for example, issuers of tickets).

We have categorized references (and their implied solutions) by actor (retailer, distributor, airline), further split into online or offline. Identifying the actors is important because the objective of each is different, and hence, the proposed models and solutions cited in the literature reflect these differences.

Fare-matching

The airline deregulation process began with the enactment in the US of the Airline Deregulation Act of 1978. Since then, deregulation has been implemented to various degrees elsewhere in the world. Deregulation has resulted in price volatility and a dramatic increase in the number of fares.

To remain competitive, airlines had to update fares more frequently in response to competitors' actions. In fare-matching, which is still a prevalent form of competitive response, airlines use monitoring tools to identify competitors' fare changes and, based on preset rules, automatically match the competitor's fares (for more information on the fare-matching practice, see Vinod, 2010).

Fare-matching is insufficient for the competitive awareness because it does not respond in real time and ignores availability and differences in schedule quality.

Strategic pricing

Unlike fare-matching, which is reactive, strategic pricing proactively files prices that would optimize an airline's expected revenue by taking into account customers' scheduling preferences (for example, number of connections, duration or fare restrictions).

Vinod *et al* (2009) and Ratliff (2015) describe how to set strategic prices by modeling a customer's scheduling preferences through the use of a discrete choice model.

While strategic pricing is an improvement over fare-matching, it still suffers from lack of responsiveness in real time and ignores the availability of the airline and its competitors.

Dynamic availability

Internet accessibility enabled the low-cost carrier (LCC) business model because fenceless fare structures were better suited for distribution over the Internet. In the absence of rules and restrictions, fare classes act as price points that the Inventory system opens and closes dynamically, depending on demand and customers' WTP. Historically, this has been called DP. As prices remain static, however, we refer to this approach as *dynamic availability*.

Gallego and van Ryzin (1994) developed the theoretical foundation for dynamic availability. They modeled demand as a homogenous (time-invariant) Poisson process. By controlling the price, the airline controlled the intensity of the demand. They derived structural properties for the lowest available price, inventory monotonicity (lowest available price increases with less inventory) and time monotonicity (lowest available price decreases with less time to departure).

Zhao and Zheng (2000) generalized these results by modeling the demand as a non-homogeneous Poisson process. They showed that inventory monotonicity continues to hold, while time monotonicity in general will not.

Westermann (2006) and Isler and D'Souza (2009) discuss how airlines that use a global distribution system (GDS) can apply dynamic availability. RMS online controls are enhanced by comparing origin and destination (O&D) marginal revenue against the sum of bid price (BP) for the traversed legs.

Dynamic availability has real-time responsiveness (hence the term *dynamic*), and is consistent with RMS. Despite its important commercial value, however, it does not consider competitors' prices, schedule quality or customer segmentation. Further, by being restricted to published price points, it also lacks pricing precision. These deficiencies should not be seen as limitations of dynamic availability but, rather, as a generic flaw of the existing RMS paradigm.

Competition in RMS

To avoid the complexity of *game theory* (see below), competition-based approaches to RMS traditionally assume that the competitor is unresponsive. By incorporating competitive pricing information into demand forecast and optimization, the idea is to improve RMS performance.

d'Huart (2010) uses PODS simulations to calculate spill-in/spill-out by comparing a competitor's current lowest available booking class at a given number of days before departure with the corresponding historical distribution of the competitor's lowest available booking classes. Assuming steady-state conditions and that the competitor does not apply user overrides, this approach tells us whether the competitor was more open (spill-out) or more closed (spill-in) than typical. PODS simulations show that this approach provides revenue gain.

Zhang and Kallesen (2008) formulate a dynamic programming model for a single leg with one competitor, for a fenceless fare structure with discrete price points. The competitor's price information is incorporated by assuming a predefined price curve that follows a stochastic process (Markov chain) generated by the competitor's price-transition matrix. They demonstrate that their model provides positive revenue gain versus traditional approaches (for example, EMSR) that ignore competition. Although these approaches were revenue-positive, their practicality is questionable, given the assumptions of steady-state conditions.

Though modeling that indirectly accounts for competitors' influences can be incorporated into offline RMS processes, this approach has no real-time component. Consequently, there is no accounting for availability or potential differences in schedule quality of the airline and its competitors.

Game theory

In game theory, the objective is to optimize revenue by considering competitors' strategies. In principle, this would require the airline to model a competitor's RMS system, calibrating the competitor's demand forecast as well as having perfect knowledge of its inventory so as to predict its strategy across the booking horizon. Obviously, this approach is impossible, and the standard solution is to analyze the existence and stability of a Nash equilibrium (*no player can gain by unilaterally changing strategy*).

Gallego *et al* (2006) show that a unique Nash equilibrium exists for games with a multinomial logit (MNL) demand model under mild and realistic assumptions associated with the utility and cost functions.

Isler and Imhof (2008) consider a single flight leg with two competitors, each with fenceless fare structures, assuming perfect knowledge of the competitor's inventory. They model customer arrival as a homogeneous Poisson process. Through simplified demand assumptions, they analyze the Nash equilibrium using dynamic programming techniques and show that the equilibrium leads to complete spiral-down for large capacities (zero bid price).

Gallego and Hu (2014) consider an O&D market with multiple competitors but are more general about the product definition. They formulate a differential game, assuming demand is deterministic. Assuming further that competitors are unresponsive, they prove the existence of a Nash equilibrium as well as conditions for its uniqueness (the MNL demand model satisfies these conditions). The equilibrium state may be characterized by a relatively simple structure (time-dependent equilibrium prices and time-independent BP representing competitors' capacity pressure). This structure allows them to conclude that this equilibrium may also be applicable to responsive competitors.

Game theory offers a more accurate model of the competitive environment and as such provides valuable insight. To our knowledge, however, it has never been used in real-world RMS applications because of its complexity.

Competitor matching

In competitor matching, the airline monitors, possibly in real time, competitors' prices and availability and responds by opening or closing booking classes based on simple market-specific business rules (for example, expected load factor or days before departure).

Lua (2007) tested these ideas by performing PODS simulations. Results were generally negative for the airline performing the matching and positive for the competitor. This was because competitor matching fails to account for the airline's own RMS strategy. In addition, the method ignores differences in schedule quality between the airline and its competitors.

Competitor matching improves the online RMS process through real time, competitive-pricing responsiveness. However, in addition to the previously mentioned limitations, competitor matching lacks consistency and pricing precision.

Agent content and pricing management

The business model for an online travel agent (OLTA) can be either an agency model or a merchant model, each of which has different implications for revenue maximization.

In the agency model, suppliers set prices and pay agent's commissions and/or override incentives, provided that certain volume targets are met. The agent may also charge customers a service fee to offset costs of doing business, such as marketing campaigns (for example, cost of Google's AdWords).

Rao and Smith (2006) and Smith *et al* (2007) describe how Travelocity, under an agency business model, applies revenue planning in response to customers' request, thus controlling

Table 1: Suppliers business needs and existing solutions within RMS

	RMS online		*RMS offline*		*Fare management*		
Business needs	*Competitor matching*	*Game theory*	*Dynamic availability*	*Competition in RMS*	*Fare-matching*	*Strategic pricing*	*Dynamic pricing*
Competitors' prices	Yes	Yes	No	Yes	Yes	Yes	Yes
Real-time responsiveness	Yes	Yes	No	No	No	No	Yes
Schedule quality	Yes	Yes	No	Yes	No	Yes	Yes
Customer segmentation	Yes	Yes	Self-selection	N/A	Fare rules	Fare rules	Yes
Price precision	No	Yes	No	N/A	Yes[a]	Yes[a]	Yes
RMS consistency	No	Yes	Yes	Yes	N/A	N/A	Yes

[a]Pricing precision is limited by the frequency of fare filing.

which itineraries to display and in which order. The objective is to maximize revenue from airline and market commissions, performance-specific incentives and service fees.

In the merchant model, the agent accesses the supplier's content at a discounted rate, leaving the agent free to mark it up. This model can be used by airlines that seek an alternative distribution channel for distressed inventory. Opaque products, in which the customer is not given specific information about the airline or schedule until after the booking, provide an example of how this model can be used.

Smith *et al* (2007) explain how Travelocity, under the merchant business model, dynamically optimizes the markup during the shopping sessions. The optimization includes both a conversion model (customer buys either opaque or non-opaque) and a share model (mix of opaque and non-opaque), considering competitive prices for the non-opaque product.

These models share many of the facets of the DP model considered in this article. However, for a travel agent, the optimization problem is simplified because the agent does not *own* the inventory and, hence, can *ignore* the RMS aspects.

To summarize our literature review, Table 1 describes how different business needs are met by each solution. These needs are categorized by:

- *Competitors' prices* – to use competitors' price information, online/offline
- *Real-time responsiveness* – to respond in real time to the context of the session
- *Schedule quality* – to account for customer schedule preferences
- *Customer segmentation* – to price differentiate by customer segment
- *Price precision* – to price in between published price points (continuous prices)
- *RMS consistency* – to account for RMS recommendations

DP FUNDAMENTALS

DP is a new and innovative concept that, to our knowledge, has not been implemented in the airline industry. Even the definition of DP is not fully established, as different definitions exist. This section will define what we mean by DP and compare and contrast it with existing RMS terminology. A summary is provided in Table 2.

RMS objective

The first RMS, based on leg control, emerged in the 1980s. The objective was to maximize revenue from each flight leg separately. This required demand forecasts as well as optimization at the leg level. In the 1990s, O&D RMS started to emerge. In these systems, the objective was to maximize revenue for the entire network, considering all future revenue streams from different O&D traffic flows. As a result, this required demand forecasting at the level

Table 2: Comparison of O&D RMS with DP

	O&D RMS	*DP*
Objective	• Maximize revenue, across the booking horizon	• Maximize contribution within the session
Demand assumptions	• O&D traffic flows are assumed independent[a] • Demand forecast are averages over shopping sessions	• O&D traffic flows are assumed dependent • Demand forecast are session specific, depending on the context
Input data by O&D flow	• Volume forecasts • Sell-up/buy-down • Yields	• Choice model parameters • All airline offers (schedules, prices …) • Bid prices from O&D RMS

[a]O&D dependency possible, but not normal practice.

of O&D traffic flow, along with network-optimization algorithms.

RMS demand assumptions

In RMS, the prevailing assumption, which we will take, is to consider demand for each O&D traffic flow independent of one another. Recently, however, Vulcano *et al* (2012) proposed demand models that overcome this limitation.

In the RMS view, demand forecasts are averages over multiple shopping sessions across customer segments and competitive situations.

RMS optimization

It is useful to review the optimization problem for a single leg from the RMS perspective (see Talluri and van Ryzin, 2004; Fiig *et al*, 2010), as we will expand on this when discussing DP. The optimization problem for a single leg with no overbooking and over a finite booking horizon $[0;T]$ can be formulated for an arbitrary demand model using dynamic programming. We proceed by considering the revenue to go, $J_t(x)$ at time t and remaining capacity x. The boundary conditions on $J_t(x)$ require that the revenue to go is 0 at the end of the booking horizon $J_T(x) = 0$ and when there is no remaining capacity $J_t(0) = 0$.

Without loss of generality, we have chosen to illustrate the concepts using a fenceless fare structure, with price $p_i > 0$, in decreasing order $i = 1, \ldots, n$. The objective of optimization is to find the fare p_i that maximizes revenue to go. Let the choice probability (quantity) $Q_t(p_i)$ and total revenue $TR_t(p_i) = p_i Q_t(p_i)$ be expressed as a function of time and price only, clearly indicating that we are not considering individual customers, but rather averages over all shopping sessions.

The Bellman recursion equation for the optimal revenue to go is then given by:

$$J_{t-1}(x) = J_t(x) + \lambda_t \max_{p_i}[\mathrm{TR}_t(p_i) - Q_t(p_i)\mathrm{BP}_t(x)].$$

Here, $\lambda_t \ll 1$ is the volume component of the demand forecast at the given departure date and time step, and the bid price is defined as $\mathrm{BP}_t(x) = J_t(x) - J_t(x-1)$. To maximize the objective function, we have to maximize the contribution $c_t(p_i) = \mathrm{TR}_t(p_i) - Q_t(p_i)\mathrm{BP}_t(x)$ in each time step. In the special case of static and deterministic demand, this corresponds with the characterization of equilibrium prices considered in Gallego and Hu (2014).

Maximization of the contribution is equivalent to the standard acceptance criterion used in online RMS: $p \geqslant \mathrm{BP}_t(x) + \mathrm{FM}_t(p)$. To see this, note that maximization is obtained by setting $\Delta c_t(p)/\Delta Q = 0 \Rightarrow \mathrm{MR}_t(p) - \mathrm{BP}_t(x) = 0$, where $\mathrm{MR}_t(p) = \Delta\mathrm{TR}_t(p)/\Delta Q$ is the marginal revenue.

Finally, substituting the fare modifier $\mathrm{FM}_t(p) = p - \mathrm{MR}_t(p)$ into the equation provides the RMS criterium. The criterium is applicable to both leg-based RMS and O&D-based RMS (for additional details, see Fiig *et al*, 2010).

RMS input data

The input data to O&D RMS is defined by the requirements of network optimization. We need a valuation and a demand forecast at the level of O&D traffic flows.

The valuation, or *yield*, is the airline's net revenue after commission, taxes and proration. This is obtained from flown coupons (revenue accounting) and published fares.

The O&D demand forecast is obtained from the RMS forecasting system, which estimates forecast parameters based on historic observations (that is, bookings, fares and availabilities) and apply these parameters to produce a demand forecast for future departure dates (for additional details, see Fiig *et al*, 2014).

DP definition

We define DP as *dynamic calculation of the optimal price, taking into account the airline's strategy, customer-specific information, and real-time alternative offerings*.

DP demand assumptions

In contrast to an averaged view of demand, DP considers individual customer characteristics (for example, segmentation) and the actual choice sets available to customers during shopping sessions. For a customer who makes travel decisions, searches or purchases using an OLTA, he or she must prioritize among potentially hundreds of itineraries, with different prices and product characteristics across multiple airlines.

To model an individual customer's decision-making process, discrete choice models have been found to be superior to traditional forecasting methodologies (Garrow, 2010). To estimate the discrete choice models' parameters, however, we need information not only about the choice selected but also about the alternatives not chosen. The fact that this information normally is not readily available has prohibited the use of discrete choice models. However, technological advancements and radical growth in online sales have made choice-related information more accessible.

A wide range of choice models is available (see Garrow, 2010 for an extensive review). In practice, the models that have attracted the most attention are MNL and nested logit models. These models lead to a closed-form expression for the choice probability, which greatly simplifies both interpretation and computation.

Applying the notation common in choice models, we define the choice probability $Q_t(i|p_i, \boldsymbol{x}_{ni}, \boldsymbol{\beta})$ for an individual customer n having a choice set $i \in C_n$, where p_i is the price, $\boldsymbol{x}_{ni}$ is the vector of non-price attributes associated with individual n and alternative i and parameter estimates $\boldsymbol{\beta}$ (time-dependence through the parameters).

DP objective

While RMS has a holistic objective – to maximize revenue for the entire network considering all future revenue streams from different O&D traffic flows – DP has a more modest objective: to calculate the optimal price for a given shopping session.

In analogy with RMS, DP seeks to maximize the contribution c_t in each time step. Unlike RMS, however, DP takes the BP as an input (from RMS) and, therefore, the volume component of the demand forecast λ_t becomes irrelevant.

This simplification allows us to expand the choice probability beyond the limitations of RMS (*independence* and *averaging*), so as to consider choice models that incorporate customer-specific identification and segmentation and real-time alternative competitive offerings.

To illustrate this, let us denote the airline that implements DP as the *DP airline*. Then, for customer n, we can calculate the contribution for the alternative $i \in C_n$, $c_t(i) = \mathrm{TR}_t(i) - \mathrm{BP}_t(x) \times Q_t(i|p_i, \boldsymbol{x}_{ni}, \boldsymbol{\beta})$. Therefore, the objective for DP is to determine the prices that maximize the sum of contributions of the alternatives across the DP airline's subset $S_{AL} \subseteq C_n$:

$$\max_{p_i} \sum_{i \in S_{AL}} c_t(i)$$

It should be noted that DP and RMS's contribution become identical if we restrict DP to the underlying assumptions of RMS (independence among alternatives, average choice probability and selection among published price points only).

DP input data

To model session-specific customer choices, input data should include the choice model parameters (by customer segment) and the choice set available to the customer. This comprises all of the product and price characteristics (for example, schedules, out-of-pocket prices, product quality) of the available offers.

DP example

We conclude this section with an example that illustrates the price and contribution impacts of DP. Let us focus on a single flight leg, with two competing airlines, AL1 and AL2, each with a fenceless fare structure. For simplicity, we ignore the influence of taxes and other surcharges. In this example, f denotes published (fixed) fares, while p_1, p_2 denote the lowest available (continuous) price for AL1 and AL2, respectively.

Both airlines offer three fares: $f = \{75, 100, 125\}$. Assume that purchase probability for AL1 follows a MNL model: $Q_1(p_1, p_2) = e^{u_1} / (\sum_j e^{u_j} + e^{u_{\mathrm{NP}}})$, where: $u_1 = -1.5 p_1/p_0 - (p_1/p_2)$; u_{NP} is the utility associated with the 'no purchase' option. Analogously, the formulas for AL2 are obtained by interchanging indices, $1 \leftrightarrow 2$. In our example, we set $p_0 = 50$, $u_{\mathrm{NP}} = -4$; the combined share of AL1 and AL2 is approximately 66 per cent when $p_1 = p_2 = 100$.

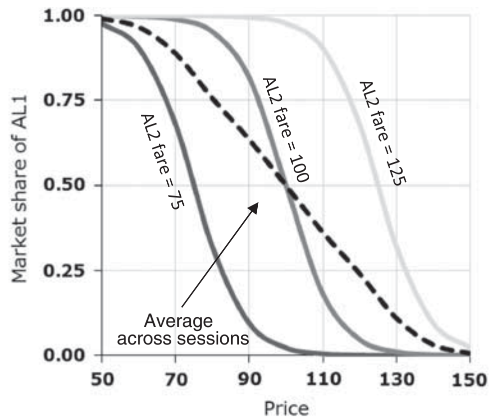

Figure 2: DP example. DP view: Market share versus price is given at the session level, for fixed competitor (OA) prices. Illustrated for AL2 price = 75, 100, 125 (solid lines). RM view: The consolidated RM view of the market share versus price (dashed line) shows lower price sensitivity.

Figure 2 illustrates the market share for AL1 versus p_1 for each possible value of p_2 (solid curves). As the RMS is based on a consolidation of booking history across *multiple* sessions, it tends to see a different averaged relationship. This is illustrated by the dashed curve. The apparent impact of price (or other non-price factors) tends to be diminished. This difference can have profound consequences.

We illustrate the impact of DP by considering the performance of AL1 in three scenarios:

(1) AL1 and AL2 both employ RMS using static filed fares.
(2) AL1 uses DP while AL2 uses RMS.
(3) AL1 and AL2 both use DP.

In each case, we measure AL1's performance in terms of expected contribution: $c_1 = Q_1 \times (p_1 - \mathrm{BP}_1)$, where BP_1 is the bid price.

Note that our bid price includes the expected cost of displacing future demand and also actual incremental costs. We include an incremental cost of 25 and consider bid prices above this level.

Case 1: AL1 RMS, AL2 RMS

In each session, AL1 can make only one fare available, and thus has to select one of $f_1 = \{75, 100, 125\}$ across all sessions. The best

Table 3a: RMS versus RMS

RMS versus RMS	*AL2: Low BP*		*AL2: Medium BP*		*AL2: High BP*		*AL1: Average*	
	p_1	c_1	p_1	c_1	p_1	c_1	p_1	c_1
AL1: Low BP	75.00	16.58	75.00	25.13	75.00	29.42	75.00	23.71
AL1: Medium BP	100.00	6.22	100.00	12.83	100.00	18.16	100.00	12.40
AL1: High BP	125.00	1.24	125.00	3.20	125.00	5.47	125.00	3.30
Average							100.00	13.14

Table 3b: DP versus RMS

DP versus RMS	*AL2: Low BP*		*AL2: Medium BP*		*AL2: High BP*		*AL1: Average*	
	p_1	c_1	p_1	c_1	p_1	c_1	p_1	c_1
AL1: Low BP	69.67	17.02	81.06	25.66	90.22	32.27	80.32	24.98
AL1: Medium BP	88.66	7.26	96.89	13.30	104.39	18.63	96.65	13.06
AL1: High BP	125.20	1.46	129.16	3.61	133.46	6.19	129.28	3.75
Average							102.08	13.93

Table 3c: DP versus DP

DP versus DP	*AL2: Low BP*		*AL2: Medium BP*		*AL2: High BP*		*AL1: Average*	
	p_1	c_1	p_1	c_1	p_1	c_1	p_1	c_1
AL1: Low BP	64.71	13.17	75.98	21.82	90.59	32.48	77.09	22.49
AL1: Medium BP	88.74	7.27	95.13	11.99	105.40	19.29	96.42	12.85
AL1: High BP	127.45	2.57	129.97	4.06	134.92	7.04	130.78	4.55
Average							101.43	13.30

price RMS can make available, p_1^{RMS}, depends only on BP_1 in a piece-wise linear fashion (breaks determined from the contribution maximization condition; see 'Dynamic Pricing Fundamentals').[1]

We denote the three ranges as low, medium, and high:

$$p_1^{\mathrm{RMS}} = \begin{cases} 75\ \mathit{for}\ 25 < \mathrm{BP}_1 \leqslant 43:\ \mathrm{BP}\ \mathit{Range} = \mathit{Low} \\ 100\ \mathit{for}\ 43 < \mathrm{BP}_1 \leqslant 79:\ \mathrm{BP}\ \mathit{Range} = \mathit{Medium} \\ 125\ \mathit{for}\ 79 < \mathrm{BP}_1 \leqslant 100:\ \mathrm{BP}\ \mathit{Range} = \mathit{High} \end{cases}$$

For this example, we assume that each bid price range is equally likely and that bid prices are uniformly distributed within each range. Further, we assume that AL2 has the same bid price distribution and that it is independent of AL1's bid price distribution. Table 3a shows how p_1^{RMS} *and* c_1^{RMS} vary with BP_1. Note that p_1^{RMS} does not depend on BP_2, while c_1^{RMS} does. Averaging across the nine equally likely scenarios yields an average price of 100 and contribution of 13.14.

Case 2: AL1 DP, AL2 RMS

In this case, DP sets p_1^{DP} knowing p_2. Table 3b shows the results. For example, when BP_1 *and* BP_2 are both in the low range, p_1^{DP} drops from 75 (Case 1, above) to 69.67; contribution increases from 16.58 to 17.02. Averaging across these scenarios yields an average price of 102.08 and contribution of 13.93, a 6 per cent increase in contribution versus RMS.

DP improved performance in two ways. First, DP sets the available price knowing the competitor's actual price. Second, DP is not restricted to the initial three fares but can select prices continuously. If we restrict DP only to the original fares, the increase in contribution is 4 per cent.

Case 3: AL1 DP, AL2 DP

We assume that AL1 and AL2 use identical DP processes and have perfect knowledge of each other's bid prices, costs, and strategies.

In this competitive game, we will assume that session prices will converge to a Nash equilibrium. The Nash equilibrium $\boldsymbol{p}^{\mathrm{DPDP}} = (p_1^{\mathrm{DPDP}}, p_2^{\mathrm{DPDP}})$ can be found from stationary conditions $\partial c_1(\boldsymbol{p}^{\mathrm{DPDP}})/\partial p_1 = \partial c_2(\boldsymbol{p}^{\mathrm{DPDP}})/\partial p_2 = 0$. While there is no analytic solution to find the equilibrium in our example, Gallego *et al* (2006) describe a 'best-response' approach to find these solutions and show a linear convergence rate. We use a similar approach within the session to find $\boldsymbol{p}^{\mathrm{DPDP}}$.The resulting prices and contributions depend on the bid prices for both airlines. Table 3c shows AL1 across the bid price ranges. In this case, prices and contributions are reduced when demand and bid prices are low. Performance is reduced below that of RMS when demand is low. When demand and bid prices are high, two-player DP performance exceeds that for RMS versus RMS (Case 1) and DP versus RMS (Case 2). Performance is mixed in the medium bid price range. Averaging across all cases shows that DP performs slightly better than RMS (1 per cent).

Figure 3 displays price and relative contribution gain (to RMS) versus BP_1 for AL1 in each scenario. The RMS price, p_1^{RMS}, is step-wise constant while p_1^{DP} increases smoothly. The strongest contribution gains of DP versus RMS and DP versus DP, relative to RMS versus RMS, are seen near the edges of the bid price ranges. This is because at the bid-price edges RMS price discontinuity provides lowest efficiency. The DP versus DP contribution gain can be negative with respect to RMS when demand and BP_1 are low.

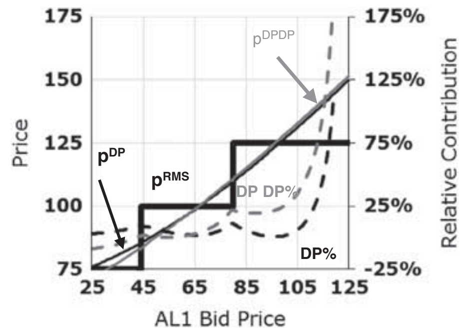

Figure 3: Optimal prices (primary vertical axis: solid lines) for AL1 is shown for three scenarios: RMS versus RMS, DP versus RMS, DP versus DP. Relative contribution (secondary vertical axis: dashed lines) for AL1 is shown for the same scenarios.

We have used a very simple example to illustrate DP behavior and performance. While the results seem positive directionally, we should not read too much into the details. Our example will tend to exacerbate the impact of DP because of unrealistic assumptions: head-to-head competition in a single market; perfect knowledge of the competitor's bid prices, cost structure and strategies; and a market share model that ignores non-price factors.

However, even with these assumptions, we see no indication of a price spiral, up or down, to unrealistic levels.

RULES-BASED DP

As discussed in the introduction, the purpose of our article has been to validate the DP concept and estimate the potential revenue gain. Our approach, therefore, has been to simplify the implementation of the DP engine, by considering a rules-based approach.

Figure 4 depicts a competitive scenario: Airline 1 (AL1), a DP-airline versus Airline 2 (AL2), an RMS-only airline. We consider a customer request for a given O&D city pair, and we assume that this request can be satisfied by both airlines that proposed different product and price offerings in terms of schedule quality, time of day, fare restrictions and booking classes. The graph on the left illustrates fares for the DP

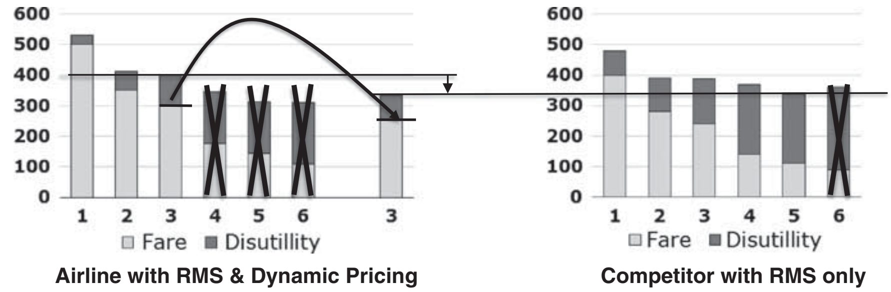

Figure 4: Scoring methodology in rule-based DP.

airline with booking classes 1–6, in decreasing fare order. The non-price disutility costs from schedule quality, fare restrictions or other factors are added on top of each bar. The aggregate bars represent the customer's total generalized cost – which we term the *perceived price*. Classes 4–6 are closed by the RMS, leaving only classes 1–3 available. On the right, the competitor's booking classes 1–6 are shown. The competitor's RMS has closed down Class 6. Assuming that the alternatives constitute the complete choice set, the customer (being rational) would likely select the lowest available perceived price within his WTP. In this case, in the absence of DP, the customer would select the competitor's Class 5, assuming that his WTP is above the fare level for Class 5.

With rules-based DP, the DP airline modifies prices according to the competition, the identified travel purpose and schedule quality, using the following rules:

- For each competitor's itineraries, the estimated perceived price (price+sum of disutility costs) of the lowest available fare is computed. The itinerary with the lowest estimated perceived price is defined as the most attractive competitor offer.
- Similarly, the most attractive offer for the DP airline is identified.
- A score based on the perceived price ratio is computed by comparing the attractiveness of the DP airlines most attractive offer *vis-à-vis* that of the competitor (details of the score computation are proprietary).
- The price is adjusted, according to the computed score and pricing rules. These rules define ranges of score values together with the associated price modifications.
- Price adjustment is capped as follows: The adjusted price of a class is bounded by the fare levels of adjacent classes. For the topmost class, there is no boundary. The adjusted price cannot fall below the bid price of the itinerary, because this would result in a negative contribution (see 'Dynamic Pricing Fundamentals').

In this example, the DP airline that applies these rules would adjust Class 3 fare to match the competitor's offering, though in practice the exact value would depend on the detailed rules and bid price.

SIMULATION ANALYSIS

As part of this study, Amadeus contracted with PRL LLC[2] to perform simulation studies on DP. PRL has noted the following caveats with respect to these results: 'PODS results are specific to the sets of network, competitive environments, and models assumptions of these particular studies, and while indicative, may change under other input conditions'. We also provide additional simulation results from ABE, an internal Amadeus simulator. In a completely different simulation, ABE confirmed the revenue gain obtained through DP.

PODS and its various components have been widely described in the literature (see, for example, Hopperstad, 1997 or theses and dissertations on MIT's DSpace@MIT Website).

The central concept in PODS is that RMS models have demand forecasts but no 'inside'

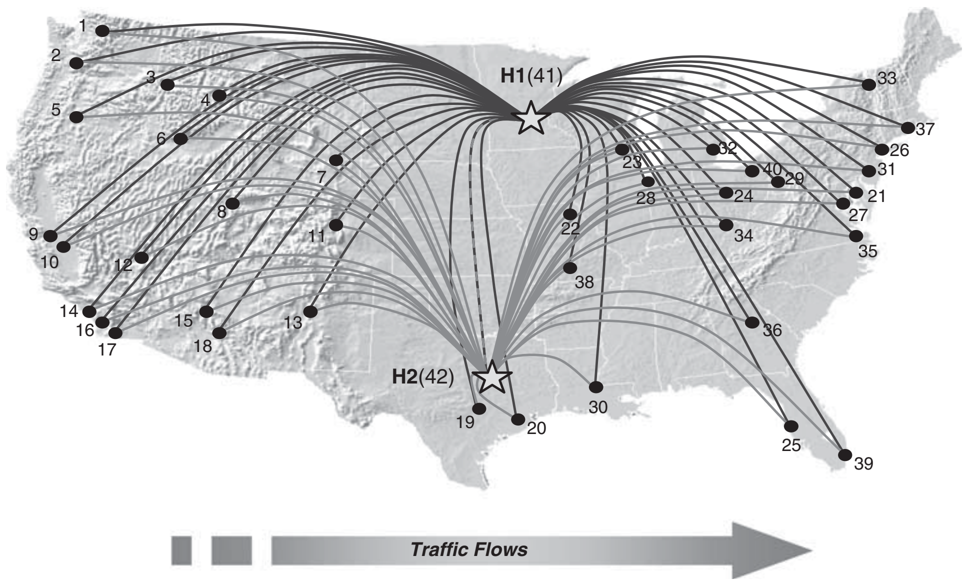

Figure 5: PODS Network D.

information about the simulators' passenger-choice model.

Each simulated passenger chooses among available alternatives (that is, O&D traffic flow and product/fare combinations) with a fare below the passenger's randomly drawn maximum WTP. If more than one alternative exists within the WTP range, the passenger chooses the alternative with the lowest perceived price (see previous section).

Figure 5 shows the geography of Network D, one of the standard PODS networks. In Network D, there are two airlines (AL1 and AL2), each with their own hubs. The traffic flow is from 20 western cities via the hubs and to 20 eastern cities. Each airline operates three connecting banks per day with 126 legs in total, which produce 1446 traffic flows in 482 city pairs (markets). The fare structure is semirestricted, with six booking classes.

In the PODS simulations, DP is implemented using the rules-based logic explained in the previous section. We assume that DP has perfect knowledge of the mean disutility parameters for both leisure and business passengers, which assumes a highly accurate and unbiased estimator on the part of the DP airline. In the real world, such levels of estimation accuracy could well be difficult to achieve.

Scenarios

We measured DP performance under various competitive conditions and supported the results through sensitivity analyses. An overview of the scenarios is provided below. Summary statistics can be found for all simulated cases in Table 4. For selected scenarios, we also provide graphical representations with additional details.

Case 1	DP versus leg RMS competitor
Case 2	DP versus O&D RMS competitor
Case 3	DP versus DP competitor
Case 4	Impact of simulation set up
Case 5	Impact of constraints on DP
Case 6	Impact of bias on DP

Case 1: DP versus leg RMS competitor

The base case in this simulation was AL1, which applies O&D RMS, versus AL2, which applies AT90, a less-sophisticated leg RMS corresponding to an LCC. AT90 adjusts the booking limits as an accordion to obtain a constant load factor. It is simplistic and entirely empirical, but has proven to be a very robust RMS.

In Figure 6 we display the impact of AL1 applying DP. Panel (a) shows the revenue relative to the baseline (index). AL1 obtained a revenue gain of 6.78 per cent, while the

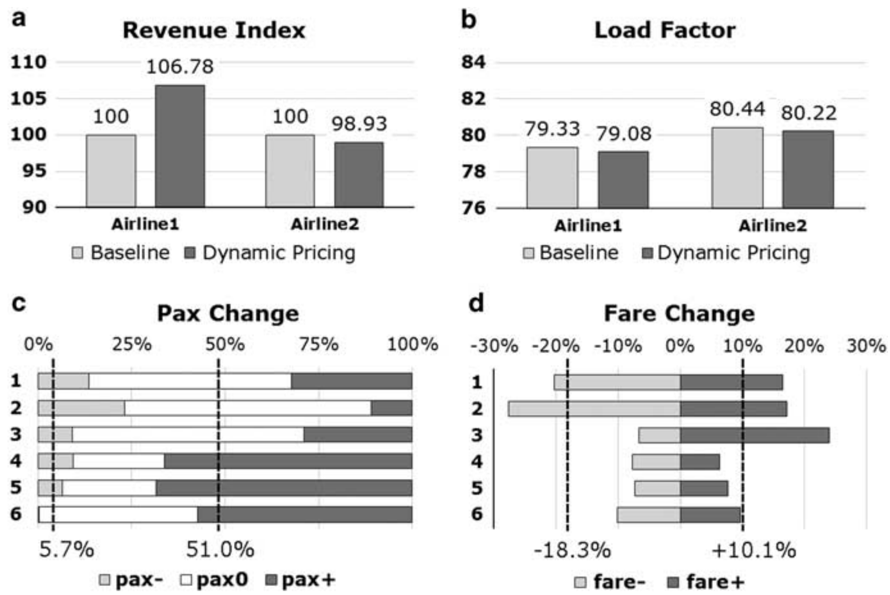

Figure 6: Impact of AL1 applies DP. Scenario O&D RMS versus AT90.
(**a**) Revenue indices; (**b**) Load Factors; (**c**) Pax distributions by booking class, showing decreases, increases or no changes; (**d**) Price change distribution by booking class for pax impacted by DP. Dashed lines shows pax weighed averages across all booking classes.

competitor, AL2, lost 1.07 per cent. The effect of DP on the load factor, Panel (b), was marginal (AL1, −0.25pp; AL2, −0.22pp).

Panel (c) displays how passenger volumes were affected by DP, categorized by price decreases (pax−), increases (pax+) or no change (pax0). The average percentage volumes affected by price increases or decreases (the dashed lines) show that AL1's revenue gain came primarily from prices increases.

Panel (d) shows the magnitude of the price changes (relative to the average fare of each booking class) for passengers affected by DP. The average magnitude is marked by dashed lines; on average, this magnitude ranged from [−18.3 per cent; +10.1 per cent] across booking classes.

It should be noted that the revenue gain from applying DP, 6.78 per cent (AL1, Panel (a)), is so substantial that it corresponds to the accumulated benefits of RMS over its entire 40-year development. Hence, independent verification is warranted. This will be provided below using the ABE simulator in a completely different set up (Case 5c).

Case 2: DP versus O&D RM competitor

This simulation investigated the impact of the competitor having a more advanced O&D RMS.

The base case, therefore, was changed to a symmetric case in which both airlines apply O&D RMS. We observed a significant but marginally lower revenue gain for the DP airline, 5.83 per cent. The competitor experienced a revenue loss of 0.88 per cent, which was slightly less than its performance in Case 1.

In this case, the load factor (LF) decreased marginally for AL1 while the competitor increased LF by 0.93pp. AL2's performance improved because its advanced O&D RMS accepted more passengers spilled by AL1 than if it were controlled by the leg-based RMS, AT90, which maintains a constant load factor.

Case 3: DP versus DP competitor

As was the case when RMS was introduced in the airline industry, it could be feared that DP is a zero-sum game – such that gains for the DP airline come at the expense of a competitor or

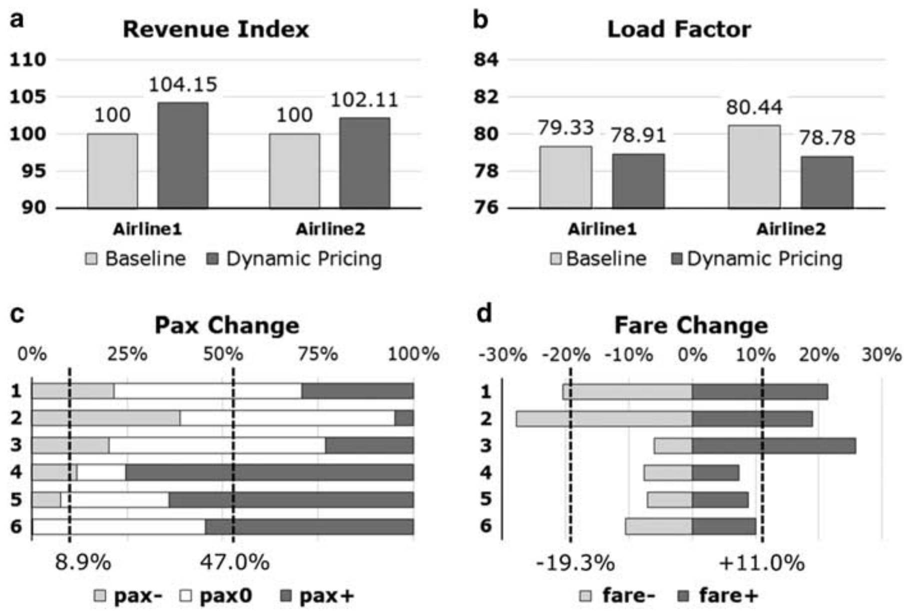

Figure 7: Impact of both AL1 and AL2 apply DP. Scenario O&D RMS versus AT90.
(**a**) Revenue indices; (**b**) Load Factors; (**c**) Pax distributions by booking class, showing decreases, increases or no changes; (**d**) Price change distribution by booking class for pax impacted by DP. Dashed lines shows pax weighed averages across all booking classes.

that DP is inherently unstable, causing prices to spiral up or down (this was the reality when traditional RMS was applied in simplified fare structures; see Cooper *et al*, 2006). However, this simulation demonstrated that DP is well behaved.

The competitive DP scenario assumes that both airlines have identical (and complete) knowledge of the other players' products and prices before they simultaneously apply DP. In this scenario, each airline does not explicitly anticipate the other airline's DP actions in its strategy. This could be viewed as unnecessarily restrictive, but in reality the customer is faced with hundreds of offerings from many competing airlines – some of which may or may not apply DP. In trying to anticipate the competitors' response, the DP airline is faced with incomplete information – not only about customer segmentation and preferences but also about the competitors' bid prices and costs. Further, the DP airline will have difficulty monitoring the competitor's responses, as DP is activated only for selected sessions and may not be reproducible because the responses depend on the real-time context.

For this simulation, results of the base case O&D RMS versus leg RMS, when both airlines apply DP, are displayed in Figure 7. In Panel (a), both airlines obtained a significant revenue gain (AL1, 4.15 per cent; AL2, 2.11 per cent). The highest gain was obtained for the airline with advanced RMS. If we look at the industry as a whole (AL1+AL2), the key statistics can be found in Table 4: Industry revenues increased by 3.17 per cent when both airlines applied DP, close to the industry gain of 3.01 per cent when only AL1 applied DP.

Hence, we can think of DP as extracting revenue from the market by reducing the customer surplus. The gain is spread among airlines that utilize DP, giving airlines an incentive to be a first mover.

Regarding loads, Panel (b) indicates that when both airlines applied DP, LF for AL1 decreased by 0.42pp, while LF for AL2 decreased by 1.66pp, with the difference attributable to the fact that DP primarily increased prices. Panels (c) and (d) show the distribution of price changes for AL1. Overall, we observed

Table 4: Summary statistics for: base case, comparison case, DP volume distribution, and price change distribution

Base case							*Comparison case*								
Case	*Simulation*	*AL1 RMS*	*AL2 RMS*	*Scenario*	*AL1 LF*	*AL2 LF*	*AL1 revenue index*	*AL2 revenue index*	*Industry index*	*AL1 LF*	*AL2 LF*	*AL1 average volume pct + (%)*	*AL1 average volume pct − (%)*	*AL1 average price pct + (%)*	*AL1 average price pct − (%)*
1	PODS	O&D	Leg	AL1→DP	79.33	80.44	106.78	98.93	103.01	79.08	80.22	51.0	5.7	10.1	−18.3
2	PODS	O&D	O&D	AL1→DP	78.80	76.28	105.83	99.12	102.51	78.58	77.21	45.1	6.1	10.1	−16.8
3a	PODS	O&D	Leg	AL1,AL2→DP	79.33	80.44	104.15	102.11	103.17	78.91	78.78	47.0	8.9	11.0	−19.3
3b	PODS	O&D	O&D	AL1,AL2→DP	78.80	76.28	102.14	103.59	102.85	78.96	77.06	44.3	8.0	11.1	−19.6
4	ABE	Leg	Leg	0% bias	81.69	78.60	105.94	99.97	102.80	80.25	78.56	24.2	1.8	12.9	−8.2
5a	PODS	O&D	Leg	AL1→DP cap	79.33	80.44	102.08	100.18	101.17	78.92	80.55	59.6	2.9	3.9	−3.3
5b	PODS	O&D	Leg	AL1→DP+	79.33	80.44	103.39	102.08	102.76	77.85	80.66	54.4	0.0	9.8	0.0
5c	PODS	O&D	Leg	AL1→DP−	79.33	80.44	103.80	97.57	100.81	79.82	80.29	0.0	6.5	0.0	−19.2
6a	PODS	O&D	Leg	80% segm. acc.	79.33	80.44	106.07	99.09	102.72	78.93	80.21	49.7	5.9	10.7	−17.0
6a	PODS	O&D	Leg	60% segm. acc.	79.33	80.44	105.37	99.46	102.53	78.37	80.31	48.6	6.2	11.3	−15.8
6b	PODS	O&D	Leg	Cached avl.	79.33	80.44	106.68	98.93	102.96	79.15	80.19	50.2	6.0	10.0	−18.1
6c	ABE	Leg	Leg	−40% bias	81.69	78.60	105.19	99.93	102.42	80.39	78.48	25.0	1.7	11.6	−8.2
6c	ABE	Leg	Leg	−20% bias	81.69	78.60	105.49	99.96	102.58	80.36	78.53	25.0	1.8	12.0	−8.4
6c	ABE	Leg	Leg	+20% bias	81.69	78.60	105.88	99.83	102.70	80.28	78.67	23.0	1.8	13.5	−8.3
6c	ABE	Leg	Leg	+40% bias	81.69	78.60	105.88	99.81	102.69	80.28	78.61	22.9	2.9	13.5	−6.8

Note: The revenue indices for both base case and the industry are 100 (not shown).

price increases and decreases of 47.0 and 8.9 per cent, respectively.

We also verified the DP versus DP competitive scenario for the symmetric case of O&D RMS versus O&D RMS. The results were consistent with the results we obtained in Case 3.

Case 4: Impact of simulator set up

In Case 1, we observed a substantial revenue benefit from DP. To obtain independent verification of this result, we employed the ABE simulator in an entirely different network set up.

The ABE simulator shares simulation concepts, RM models, and architecture with PODS. Differences exist in the simulations' set ups; whereas PODS applies synthetic networks, schedules, fare structures, demand levels and other related factors, ABE aims to provide more realism into the simulation by building and calibrating on real airline networks, real schedules and fare structures. The network used in ABE corresponds to a subnetwork of PODS Network D (Figure 4) (with fewer spoke cities). The fare structure for ABE is similar – semirestricted with six fare classes but not identical. Our baseline was leg RMS versus leg RMS, where AL1 applies EMSR with hybrid forecasting and fare adjustment (HFFA, see Fiig *et al*, 2010 for additional details), while AL2, as before, applies AT90.

In this simulation, the DP revenue benefit was 5.94 per cent (Table 4). This is the same order of magnitude as obtained in Case 1 from PODS (6.78 per cent).

Case 5: Impact of constraints on DP

The average price adjustments in Case 1 were substantial [−18.3 per cent; +10.1 per cent], and in practice, airlines would monitor and cap the proposed price adjustments. By capping we also sought to investigate the cause of DP's revenue gain in Case 1. Did the gain come predominantly from price increases or decreases, or was it evenly distributed?

We repeated the base-case scenario (O&D RMS versus leg RMS), but this time with a constrained DP. We considered three constraints separately:

Case 5a	DP limited to fine-tuning the price within the range of [–4 per cent; +4 per cent]
Case 5b	DP applies only price increases
Case 5c	DP applies only price decreases

Case 5a: For DP to be well behaved, revenue gains should be rooted evenly across price adjustments rather than rooted in few large adjustments. Hence, revenue gains should scale with the size of the capping. To test this, we applied capping, reducing the DP price range from [−18.3 per cent; +10.1 per cent] to [−4 per cent; +4 per cent] about a factor of 4 (pax weighted). The reduction of AL1's revenue gain was of the same order or magnitude (from 6.78 to 2.08 per cent). As before, the load factors for AL1 and AL2 were practically unaffected by DP.

Cases 5b/5c: In these series of scenarios, we investigated the root cause of the DP gain. From an RMS (inward-looking) perspective, we could naïvely assume that DP would slightly adjust the optimal price point on the price – demand curve. However, this is not what happened. Unlike RMS, DP is outward looking.

The revenue gains when prices increased, came from the DP airline's superior itineraries (that is, higher schedule quality) that could sustain a price increase. The revenue gains when prices decreased, came from the DP airlines' inferior itineraries (that is, lower schedule quality) that were priced competitively.

With increasing prices only, the revenue gain for AL1 was reduced from 6.78 to 3.39 per cent. AL2 also benefitted from AL1's application of DP, because AL1 increased prices unilaterally. The price increases caused AL1 to lose −1.48pp in LF, while the competitor was almost unaffected.

With decreasing prices only, the revenue gain for AL1 was reduced from 6.78 to 3.80 per cent. The gain came at the expense of AL2, which lost 2.43 per cent. This is the first time we observed that DP could have adverse effects on the competitor. Interestingly, despite DP that allowed prices only to decrease, AL1's yield increased because AL1 was intelligent about which customer to accept – capturing high-revenue passengers from AL2. This illustrates how the underlying mechanism of DP is fundamentally different from RMS thinking.

DP brought as much revenue through price increases (3.39 per cent) as with price decreases (3.80 per cent), but the number of affected passengers was noticeably different: 54.4 per cent of passengers' prices increased, while 6.5 per cent of passengers' prices decreased.

Case 6: Impact of bias on DP

It is important that DP displays robustness in the face of input data errors that, unavoidably, exist in practice. Therefore, we conducted three series of sensitivity analyses. The base case considered was, as before, O&D RMS versus leg RMS for the two first series (Cases 6a and 6b), while Case 6c used the ABE simulator in the different network set up.

Case 6a Impact of customer segmentation bias
Case 6b Impact of competitor availability bias
Case 6c Impact of disutility estimation bias

Case 6a: In the previous simulations, we assumed perfect knowledge of segmentation (business versus leisure passengers). In this study, we investigated the impact of uncertainty in this segmentation. Simulations added noise, such that the segmentation accuracy was reduced from 100 per cent accuracy (perfect segmentation) to 80 or 60 per cent. For reference, 100 per cent accuracy provided a revenue gain of 6.78 per cent (Case 1). From this study, we observed revenue gains of 6.07 per cent for 80 per cent accuracy, and 5.37 per cent for 60 per cent accuracy. Thus, even for very high segmentations errors, the DP model remained robust. The change in load factor when segmentation accuracy was reduced was minor.

Case 6b: In the previous simulations, we assumed perfect knowledge of competitor availability in real time. In reality, this is rarely the case. Most likely, competitor availability will be supplied by third-party companies (such as INFARE or QL2) or otherwise captured during nightly batch jobs. Therefore, the availability used in DP will not be completely accurate. We investigated the impact of this inaccuracy by applying a cached availability corresponding to using 'outdated' availability obtained from the previous day.

For reference, perfect knowledge of competitor availability provided a revenue gain of 6.78 per cent (Case 1). In this study, we observed a minor reduction, to 6.68 per cent, when applying cached availability. Thus, the DP methodology was robust with respect to errors in competitor availability.

Case 6c: In previous simulations, disutility costs were unbiased. The studies were performed by applying an estimation bias to the total disutility cost (from booking class restrictions and schedule quality). If we underestimate disutility costs, the price component of the perceived price increases and customers effectively become more price sensitive. On the other hand, if we overestimate disutility costs, the quality component of the perceived price increases and, hence, customers effectively become less price sensitive.

For reference, the revenue benefit for the unbiased case was 5.94 per cent (Case 4).

Underestimating the disutility costs resulted in small reductions in revenue gain: 5.49 per cent (bias −20 per cent), 5.19 per cent (bias −40 per cent). Overestimating the disutility costs reduced the revenue gain only marginally: 5.88 per cent (bias +20 per cent) and 5.88 per cent (bias +40 per cent). Consequently, rules-based DP was robust with respect to estimation bias in the disutility costs.

LIMITATIONS

The benefits of DP in this article assume that the market competitive information is both current and comprehensive. In practice, both conditions could be violated. The volume of availability requests is surging from ever-increasing look-to-book ratios and from growth in passenger travel. This has forced GDS's to pay close attention to the performance of the availability calculation and to rely on strategies, such as proxies or caches, to mitigate the load. Thus, the information used in the DP engine may not always be completely current. We studied the impact of cached availability in Case 6b showing little degradation in DP revenue gain. However, attention to the quality of availability is essential for the successful deployment of DP.

GDS-based search transactions do not always capture information from all airlines because some airlines do not (or only partially) participate in GDS channels; nor do they capture information from other modes of transportation (for example, rail or bus). Hence, the choice set may be incomplete to various degrees. Although it is possible to add content (for example, from INFARE and QL2), the impact of incomplete information may be significant in some markets.

Finally, the trend among airlines toward branded fares means that bundled products differentiate across airlines. This will complicate the comparison of the utilities for the DP engine and, if choice models do not properly reflect this complexity, there is a risk for greater error in the calculated price.

FUTURE RESEARCH

We believe that there is potential for improvement in DP performance in several areas. One such area is to improve price optimization by calculating the optimal price that maximizes contribution. More ambitiously, DP may be utilized to build a self-consistent RMS–DP system, in which RMS anticipates DP's subsequent actions when calculating the bid prices.

Another area of inquiry could be to model the interdependence of O&D traffic flows, in particular when an airline offers service between city pairs multiple times per day or with different schedule quality. Unlike traditional RMS, DP has the ability to dynamically adjust prices on all flows simultaneously.

CONCLUSION

Existing RMS do not explicitly model competition. This simplification worked well in a stable business environment, with static fares, restrictions, and schedules.

However, the current business environment – characterized by fast-changing schedules, products, fare restrictions and prices – makes this simplification inadequate, because RMS is blind to the real-time competitive situation.

Over the years, multiple attempts have been made to incorporate competitive information into RMS, with little success. We propose a new approach, DP, in which pricing is optimized in real time depending on contextual information (customer segmentation and competitor availability).

We studied the performance of DP using two simulators: PODS and ABE. We have shown that a rules-based DP engine, with no changes in existing RMS, provides significant revenue gains, in the order of ~6 per cent. When both airlines apply DP, revenue gain persists and is split among the airlines. Using sensitivity studies, we have shown that DP is remarkably robust to errors in the input data.

For airlines to realize the full revenue benefit, DP needs to be applied consistently across all sales channels and to all shopping sessions. DP will require additional data exchanges between the actors in the distribution chain.

As a consequence, all shopping platforms and downstream processes (such as ticketing, servicing, revenue accounting, RM and interline settlement) will need to take DP into account. This will require industrywide expertise and involvement.

Finally, we would like to state that Amadeus started the construction of the DP product and manages the effect of deploying DP in all sales channels.

ACKNOWLEDGEMENTS

The authors would like to thank Valerie Viale, senior manager product marketing at Amadeus and who is responsible for the DP product, for her constructive feedback. In addition the authors are also indebted to our anonymous referees for their time and exceptional valuable comments.

NOTES

1. RMS averaged view: $\overline{Q}(75) = (Q(75,75) + Q(75,100) + Q(75,125))/3 = 0, 578$; $\overline{Q}(100) = 0.322$; $\overline{Q}(125) = 0.147$. Insert into $\mathrm{MR} = \Delta\overline{\mathrm{TR}}/\Delta\overline{Q}$: $\mathrm{MR}_1(75) = (75 \cdot 0.578 - 100 \cdot 0.322)/(0.578 - 0.322) = 43.5$. $\mathrm{MR}_1(100) = (100 \cdot 0.322 - 125 \cdot 0.147)/(0.322 - 0.147) = 79.1$.
2. The PODS simulator was originally developed at the Boeing Company in the mid-1990s by Craig Hopperstad. It is currently owned by PODS Research LLC. Its research efforts have been directed, and its results validated, over the last 20 years by the airline members of the MIT/PODS Consortium.

REFERENCES

Cooper, W.L., Homem-de-Mello, T. and Kleywegt, A.J. (2006) Models of the spiral-down effect in revenue management *Operations Research* 54(5): 968–987.

d'Huart, O. (2010) A competitive approach to airline revenue management. Unpublished Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

DSpace@MIT. http://dspace.mit.edu/, accessed 29 September 2015.

Fiig, T., Härdling, R., Pölt, S. and Hopperstad, C. (2014) Demand forecasting and measuring forecast accuracy in general fare structures. *Journal of Revenue and Pricing Management* 13(6): 413–439.

Fiig, T., Isler, K., Hopperstad, C. and Belobaba, P. (2010) Optimization of mixed fare structures: Theory and applications. *Journal of Revenue and Pricing Management* 9(1/2): 152–170.

Gallego, G. and Hu, M. (2014) Dynamic pricing of perishable assets under competition. *Management Science* 60(5): 1241–1259.

Gallego, G., Huh, W.T., Kang, W. and Phillips, R. (2006) Price competition with the attraction demand model: Existence and unique equilibrium and its stability. *Manufacturing & Service Operations Management* 8(4): 359–375.

Gallego, G. and van Ryzin, G. (1994) Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science* 40(8): 999–1020.

Garrow, L. (2010) Customer Modelling. In: T. J. Hendershott (ed.) *Discrete Choice Modelling and Air Travel Demand: Theory and Applications*. Aldershot, UK: Ashgate Publishing.

Hopperstad, C. (1997) PODS Modeling Update, AGIFORS Yield Management Study Group Meeting, Montreal, Canada.

Isler, K. and D'Souza, E. (2009) GDS capabilities, OD control and dynamic pricing. *Journal of Revenue and Pricing Management* 8(2): 255–266.

Isler, K. and Imhof, H. (2008) A game theoretical model for airline revenue management and competitive pricing. *Journal of Revenue and Pricing Management* 7(4): 384–396.

Lua, W.F. (2007) Matching of lowest fare seat availability in airline revenue management systems. Unpublished Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

Rao, B. and Smith, B. (2006) Decision support in online travel retailing. *Journal of Revenue and Pricing Management* 5(1): 72–80.

Ratliff, R. (2015) Airline strategic fare optimization. Presented at AGIFORS Symposium, Washington DC.

Smith, B., Darrow, R., Elieson, J., Guenther, D., Rao, V. and Zouaoui, F. (2007) Travelocity becomes a travel retailer. *Interfaces* 37(1): 68–81.

Talluri, K. and van Ryzin, G. (2004) Revenue management under a general discrete choice model of consumer behavior. *Management Science* 50(1): 15–33.

Vinod, B. (2010) The complexities and challenges of the airline fare management process and alignment with revenue management. *Journal of Revenue and Pricing Management* 9(1/2): 137–151.

Vinod, B., Narayan, C. and Ratliff, R. (2009) Pricing decision support: Optimising fares in competitive markets. *Journal of Revenue and Pricing Management* 8(4): 295–312.

Vulcano, G., van Ryzin, G. and Ratliff, R. (2012) Estimating primary demand for substitutable products from sales transaction data. *Operations Research* 60(2): 313–334.

Westermann, D. (2006) (Realtime) dynamic pricing in an integrated revenue management and pricing environment: An approach to handling undifferentiated fare structures in low-fare markets. *Journal of Revenue and Pricing Management* 4(4): 389–405.

Zhang, D. and Kallesen, R. (2008) Incorporating competitive price information into revenue management. *Journal of Revenue and Pricing Management* 7(1): 17–26.

Zhao, W. and Zheng, Y. (2000) Optimal dynamic pricing for perishable assets with nonhomogeneous demand. *Management Science* 46(3): 375–388.