Total Revenue Optimization with the Ancillary Marginal Demand and Ancillary Marginal Revenue Transformation Heuristics

Adam Bockelie[a,*], Peter Belobaba[a]

[a]*Massachusetts Institute of Technology, International Center for Air Transportation*
77 Massachusetts Avenue, Cambridge, MA 02139

Abstract

Changes in airline business models over the last ten to fifteen years have led to a rapid growth in ancillary services and ancillary revenues. However, the development of revenue management models has not kept pace in this domain; availability controls are still designed to maximize ticket revenues, with ancillary revenues as an afterthought. A more comprehensive approach would favor booking policies that maximize total revenues.

We propose an Ancillary Choice Dynamic Program for total revenue optimization that explicitly incorporates the revenues and passenger choice impacts of ancillary services in addition to ticket revenues. We use an estimate of conditional passenger choice probabilities to compute choice and ancillary-adjusted marginal revenues for booking policies through an Ancillary Marginal Revenue transformation, and we develop an Ancillary Marginal Demand forecasting model to estimate demand volumes. We combine the AMD and AMR frameworks as total revenue optimization heuristics for existing RM optimizers, such as EMSR. We discuss implementation challenges and use the Passenger Origin-Destination Simulator (PODS) to illustrate the performance of our approach versus traditional RM models. The results suggest that AMD and AMR can increase revenue for airlines by up to 2%. Finally, we discuss the potential for our dynamic program to be used as the basis of an offer generation system, leveraging the power of New Distribution Capability.

Keywords: airline revenue management, ancillary services, total revenue optimization, passenger choice, offer generation, New Distribution Capability

1. Introduction

For decades, airlines have invested in revenue management (RM) systems to maximize the proceeds from ticket sales. Since the mid-2000s, however, airlines have been developing a secondary, ancillary, revenue stream by both unbundling their product and offering new services for sale. As the number, price, and value to airlines of ancillary services grow, so does the potential benefit of a new generation of revenue management models that attempt to maximize total revenue, not

---

[*]Corresponding author
Email address: `bockelie@mit.edu` (Adam Bockelie)

just ticket revenue. We develop a new approach to total revenue optimization with the Ancillary Choice Dynamic Program (ACDP), which explicitly incorporates ancillary revenues and the passenger choice impacts of ancillary services. We then show that the model leads to two heuristics, an Ancillary Marginal Demand transformation (AMD) and an Ancillary Marginal Revenue transformation (AMR), which together can transform existing RM models to be both ancillary-aware and choice-aware in their calculation of fare class booking limits. Finally, we discuss how emerging technologies in airline distribution will allow further extensions to our model, and describe how our work could provide a platform for an offer generation engine.

We focus specifically on optional services sold to passengers in conjunction with a particular itinerary, such as checked baggage, seating upgrades and seat assignments, inflight meals and entertainment, priority boarding, and lounge access. These services provide a significant and growing share of revenue to airlines: upgrade revenues for Delta's additional legroom Comfort+ seating section generated $125 million in 4Q 2015, and the airline expected these revenues to grow.[1] Spirit Airlines earns more than $110 million, or 5% of total operating revenue, from seat assignment fees.[2] Baggage fees provided nearly $1.5 billion for American Airlines in 2016.[3]

Most research on airline ancillary services focuses on how passengers value the services (e.g. Espino et al. (2008), Balcombe et al. (2009), and Mumbower et al. (2015)), or how the growth in services has led to changes in average fares and passenger volumes (e.g. Ancarani et al. (2009), Scotti and Dresner (2015), Brueckner et al. (2015), and Zou et al. (2017)). The economics literature contains some work studying optimal pricing/bundling decisions for generic add-on products, with a secondary focus on how consumers evaluate add-ons in conjunction with base goods (e.g. Gabaix and Laibson (2006) and Shulman and Geng (2013)). Surprisingly little research has focus specifically on how airline passengers consider ancillary services in conjunction with decisions on base ticket purchases: the first real integration is Bockelie and Belobaba (2017), who propose two behavior types: sequential and simultaneous. In the Bockelie and Belobaba framework, sequential passengers are initially myopic about ancillary charges and select an itinerary and fare class based on non-ancillary attributes, then consider ancillary services afterward. Simultaneous passengers, on the other hand, have a single unified decision process that completely integrates ancillaries with fare class and itinerary attributes.

Similarly, little work exists on integrating ancillary services into revenue management forecasting and optimization. Initial RM research focused on ticket revenue maximization and assumed independent demand streams for each fare classes. Although passengers clearly do make choices among fare products, early fare structures had few enough classes and sufficient restrictions in place to ensure each fare class was, more or less, purchased by a single segment of travelers. These models

---

[1] Delta Air Lines Earnings Call (4Q 2015)
[2] Spirit Airlines Form 10-K (2016)
[3] US DOT Form 41, Schedule P-1.2

maximized ticket revenue by flight leg (e.g. EMSR (Belobaba and Weatherford, 1996) or leg DP (Lautenbacher and Stidham, 1999)) or across networks (e.g. DAVN (Smith and Penn, 1988) or ProBP (Bratu, 1998)).

As low cost carriers have grown and implemented less restricted fare structures, this independent demand assumption has been challenged. A growing body of RM-related optimization and forecasting research has attempted to account for passenger choice amongst various fare classes. Talluri and van Ryzin (2004) propose a single-leg dynamic programing formulation that explicitly assumes a general model of passenger choice. They show that the optimal solution relies on a series of efficient sets.

Forecasting with passenger choice is more complex than with independent demand, because dependencies between classes must be considered. Hopperstad and Belobaba (2004) develop the Q-forecasting process for completely unrestricted fare structures, in which an airline forecasts total demand for a flight/market at the lowest fare class (denoted "Q"), and then uses estimated sell-up rates to partition that demand into higher-value fare classes. Boyd and Kallesen (2004) proposed separately forecasting *yieldable* demand (where passengers only purchase their preferred fare class) from *priceable* demand (where passengers purchase the least expensive fare offered, regardless of the restrictions), an approach known as hybrid forecasting (HF). The yieldable demand segment is forecasted with traditional (independent demand) standard forecasting methods; the priceable segment with a Q-forecast.

Fiig et al. (2010) and Walczak et al. (2010) developed marginal revenue transformations and marginal demand transformations for the inputs to the RM optimizer based on a model of customer choice. The transformations feed the optimizer with the expected incremental revenue or demand to be captured by opening an additional fare class, taking into account the potential for buy-down. These transformations allow any independent-demand revenue management optimization model to be adapted to account for passenger choice, enabling traditional RM systems to operate in an unrestricted or semi-restricted fare environment. The operationalized form of the marginal revenue transformation is often referred to as fare adjustment (FA) and paired with hybrid forecasting as HF/FA. Together, these extensions and models provide a mechanism to directly incorporate choice behavior into RM models.

Previous work incorporating ancillary revenue streams into revenue management models is limited. The only detailed theoretical work found in the literature is Zhuang and Li (2012). They develop a dynamic programming formulation for hotel casinos to combine room revenue with gaming revenue, equivalent to an independent demand, multiple fare class, single leg problem. The need to account for ancillary revenue streams in RM, though, has long been acknowledged (Phillips, 2005). Metters et al. (2008) reports that Harrah's, a hotel casino chain, operates its hotel revenue management system on expected total nightly contributions—the amount that an individual spends on a hotel room plus their expected gambling losses. No mathematical or numerical details are provided.

The most relevant total RM work is that of Hao (2014), who simulates the revenue performance of a total revenue optimization heuristic (which we call the optimizer increment (OI)) in a two-airline competitive environment. In Hao's simulations, passengers are assigned an expectation of ancillary "spend" based on their market and selected fare class, and airlines always collect the expected ancillary spend of every passenger. Passenger choice is not affected by the ancillary offer. Hao finds that the heuristic leads to more low-fare seat availability, which increases bookings in lower value fare classes and decreases bookings in higher value fare classes.

To summarize, although there has been recent focus on the value of ancillary services, and on the importance of accounting for passenger choice in RM models, there is no work that directly connects the two concepts. We develop a single leg RM optimization and forecasting model to address this gap, using a flexible customer choice model. Our model explicitly accounts for ancillary revenues and their impacts on passenger choice. We discuss the significant practical marketing and distribution constraints that restrict the types of offers and that airlines can sell, and show how our model can be restricted to produce booking policies that can be implemented under these conditions. We develop two heuristics that can be used to convert existing RM models to total revenue management and we show that under specific choice models our heuristics are equivalent to existing total revenue optimization methods, but that in general our models provide an additional level of specificity. We describe additional approximations and processes that may be required to operationalize our models, and finally assess their performance in detail using the Passenger Origin-Destination Simulator, finding that our models produce revenue gains across a variety of environments.

The remainder of this paper is organized as follows. Section 2 presents our model formulation, discusses practical constraints, introduces our two heuristics, and investigates their equivalence to previous approaches. Section 3 addresses challenges that may be encountered in implementing our heuristics and proposes two solution processes: one addressing the presence of inefficient booking policies, another providing a demand forecasting module. Section 4 presents our simulation results, providing an overview of the software used and the simulation environment. Section 5 provides a short introduction to New Distribution Capability and describes how our model could be used as the basis for an NDC-style offer generation engine. We conclude with a summary in Section 6 as well as thoughts on potential future work.

2. Model Formulation

We consider a single airline, single flight leg network, with multiple fare classes and multiple ancillary services. The fare classes are indexed $1, \ldots, k, \ldots, n_{FC}$ and ordered by decreasing fare; the fare for class k is f_k, so $f_1 \geq f_2 \geq f_k \geq f_{n_{FC}}$. The airline has grouped its ancillary services into purchasable combinations $0, 1, \ldots, m, \ldots, n_{COMB}$, where set 0 corresponds to the set of no ancillary

services. These combinations are formulated subject to the airline's marketing policies and goals (and ensure that passengers combine ancillary services in a sensible manner, such as a prohibiting a second checked bag without also buying a first checked bag); the set $M_k \subseteq \{0, 1, \ldots, n_{COMB}\}$ lists the combinations that are permitted in class k. Combination m of ancillary services purchased in conjunction with fare class k has price a_{km}. For modeling convenience, we assume that the fare class 0 corresponds to a decision by the passenger to not fly.

Time is discrete and counts down to departure, which occurs at $t = 0$. We assume that demand has a Poisson distribution, that time slices are small enough that there is at most one arrival per slice, and the probability of an arrival during slice t is λ_t. Capacity x is the number of unsold seats, which constrains the total number of sales. We assume that individual ancillary services have no capacity constraints and have a negligible marginal cost to the airline. We assume no cancellations or overbooking.

The airline's booking policy for each time and capacity state (t, x) is an *offer set O*. We define an offer $(k, m) \in O$ as a specific fare class k and combination $m \in M_k$ of ancillary services; a consumer purchases exactly one offer in its entirety from the offer set. The offer $(0, 0)$, which corresponds to the no-fly option, is always included in the offer set. Ancillary services are considered "optional" for a fare class if there are offers for that class in the offer set both with and without the ancillary service.

We assume that consumers make choices according to a flexible and general choice model, where the probability that a consumer chooses a particular offer is a function of the booking policy in effect and the time at which the consumer arrives in the booking process. We specify this choice model in terms of choice probabilities $P_{kmt}(O)$, which is the probability that a consumer arriving at time t chooses offer (k, m) when presented with offer set O, with $P_{kmt}(O) = 0$ if $(k, m) \notin O$. The exact structure of this choice function can vary by context; two potential models are the sequential or simultaneous behaviors described by Bockelie and Belobaba (2017).

Multiple consumer demand segments may be present, each with their own choice function. However, we assume that the airline cannot provide different booking policies to these segments, and so the probabilities $P_{kmt}(O)$ reflect a weighted average of the segments arriving at time t.

The probability that the airline sells offer (k, m) during time t, given that it has booking policy O in effect, is $\lambda_t P_{kmt}(O)$. The probability that the airline sells nothing at time t is $(1 - \lambda_t) + \lambda_t P_{00t}(O)$, which reflects that the lack of sale may be due to no arrival, or because the consumer chose not to fly. The total probability of sale TP_t for a particular booking policy is the probability that an arriving consumer purchases anything from the policy, and is:

$$TP_t(O) = \sum_{(k,m) \in O} P_{kmt}(O) \tag{1}$$

5

Likewise, the total expected revenue TR_t from an arriving customer presented with policy O is:

$$TR_t(O) = \sum_{(k,m)\in O} P_{kmt}(O)(f_k + a_{km}) \tag{2}$$

The airline selects the booking policy that maximizes the total expected revenue to come V in future time periods via the Ancillary Choice Dynamic Program (ACDP):

$$V(t,x) = \max_O \left\{ \sum_{(k,m)\in O} \lambda_t P_{kmt}(O)\left(f_k + a_{km} + V(t-1, x-1)\right) \right.$$

$$\left. + \left(\lambda_t P_{00t}(O) + 1 - \lambda_t\right) V(t-1, x) \right\} \tag{3}$$

We define a bid price function $\Delta V(t,x) = V(t,x) - V(t, x-1)$ as the marginal cost of capacity and can rewrite Equation 3 in simpler terms:

$$V(t,x) = \max_O \left\{ \sum_{(k,m)\in O} \lambda_t P_{kmt}(O)(f_k + a_{km} - \Delta V(t-1, x)) \right\} + V(t-1, x) \tag{4}$$

The airline chooses the one offer set, or booking policy, O in each time and capacity state that maximizes the total expected revenue earned from that time and capacity state, minus the cost of consumed capacity, plus the maximum expected revenue to come in future time periods. For modeling convenience, it is possible that the airline would only "offer" the no-fly option. This function is solved recursively, with the boundary conditions $V(0, x) = 0$ and $V(t, 0) = 0$: no future revenue can be earned when there is no time and/or capacity remaining.

We can redefine an offer (k, m) as a *product* j sold by the airline, with price $r_j = f_{k(j)} + a_{k(j),m(j)}$ where $k(j)$ and $m(j)$ map the product index to the fare class and ancillary combination indices. Product j has purchase probability $P_{jt}(O) = P_{k(j),m(j),t}(O)$ when the airline is selling the set of products O. We can now rewrite Equation 4 as:

$$V(t,x) = \max_O \left\{ \sum_{j\in O} \lambda_t P_{jt}(O)(r_j - \Delta V(t-1, x)) \right\} + V(t-1, x) \tag{5}$$

which recovers the (time-varying) extension of the choice-based dynamic program proposed by Talluri and van Ryzin (2004), although our use case contains the addition ancillary dimension.

Talluri and van Ryzin prove several important points about the optimal solutions to Equation 5 (in the non-time-varying case):

1. When input parameters λ_t, $P_{jt}(O)$, and r_j are known and accurate, and when demand is Poisson, the solutions to Equation 5 are optimal booking policies.
2. The optimal policies are always one of the *efficient sets*.

Because our model is mathematically equivalent, the same conclusions hold.

An efficient set, or efficient policy, is one which maximizes total expected revenue for any given total sale probability, or is part of a linear combination of policies that maximizes total expected revenue for any given total sale probability; the linear combination is, in practice, equivalent to alternating between booking policies. Talluri and van Ryzin define a policy O as *inefficient* in time slice t if there exists a set of policy weights $\alpha_t(S)$, with $\sum_{\forall S} \alpha_t(S) = 1$, such that:

$$TP_t(O) \geq \sum_{\forall S} \alpha_t(S)TP_t(S) \quad \text{and} \quad TR_t(O) < \sum_{\forall S} \alpha_t(S)TR_t(S)$$

otherwise O is efficient in time t (Talluri and van Ryzin, 2004, Section 3.1). In other words, a set is inefficient if there is a weighted combination of other sets with a smaller (or equal) sale probability and a greater total expected revenue. The no-fly-only policy (offering $(0,0)$ in our notation) is always efficient.

We can index and order the efficient sets in time t by increasing sale probability $O_{1,t}, \ldots, O_{n_{ES_t},t}$ with $TP_t(O_{i,t}) \leq TP_t(O_{i+1,t})$; the sets are therefore also ordered in terms of increasing expected revenue (Talluri and van Ryzin, 2004, Proposition 3). Fiig et al. (2010), Walczak et al. (2010), and Gallego (2013) develop marginal revenue and marginal demand transformations to convert Equation 5 into an equivalent independent demand formulation. We extend these ideas to the ancillary dimension as the Ancillary Marginal Demand transformation (AMD, Equation 6) and Ancillary Marginal Revenue transformation (AMR, Equation 7):

$$d_{i,t} = \lambda_t \left(TP_t(O_{i,t}) - TP_t(O_{i-1,t}) \right) \tag{6}$$

$$f'_{i,t} = \frac{TR_t(O_{i,t}) - TR_t(O_{i-1,t})}{TP_t(O_{i,t}) - TP_t(O_{i-1,t})} \tag{7}$$

where d_{it} is the marginal, or additional, demand that can be accommodated and f'_{it} is the marginal total revenue per unit of capacity that can be earned by moving from policy $O_{i-1,t}$ to O_{it} in time t, and we define $TR_t(O_{0,t}) = TP_t(O_{0,t}) = 0$. Fiig et al. (2010) show that these marginal unit revenues

are decreasing in i, and that the optimal policy in time t is to offer the set $O_{i_t^*,t}$ with the smallest expected revenue greater than or equal to the bid price:

$$i_t^* = \max\{i \mid f_{it}' \geq \Delta V(t, x)\}$$

2.1. Practical Constraints and Limitations

Airline marketing policies and distribution technology impose significant practical constraints on the types of offers and offer sets that airlines can sell. Traditionally, airlines have sold tickets to consumers directly (through the airline website, call center, and ticket offices) and indirectly (through travel agents and online travel retailers). Airline sales through indirect channels are subject to the technical limitations of distribution technology, and sales through all channels are potentially subject to commercial agreements with various retailers. In practice, for many airlines, the same booking policy must be in place for consumers shopping in all channels (direct and indirect).

Indirect sales are often made through a Global Distribution System (GDS), which serves as a content aggregator. Approximately 50% of bookings worldwide are made through a GDS, and therefore the structure of GDSs has a significant impact on how airlines sell travel (Taubmann, 2016). A consumer makes a shopping request to a travel retailer (such as a human travel agent, or an online travel agent like Expedia). The travel retailer then requests search results from a GDS. The GDS draws upon schedule data from a third party (Official Airline Guide, OAG), fare and pricing data from a third party (Airline Tariff Publishing Company, ATPCO), and availability data from airlines. The airline availability data lists the number of seats available in each booking class. The GDS then combines all of this information to assemble the travel options returned to the retailer (and in turn to the consumer). Figure 1 shows a schematic of this process.

Airlines are responsible for providing their schedules to OAG and fares to ATPCO, and can update the data as necessary. However, the only "real time" control that the airline has in this process is the availability response. Thus, the airline can only control products at the fare class level. In the example in Figure 1, the airline has responded that fare classes Y, B, and M are available, and that the airline sells two ancillary products, BAG1 and BAG2. The airline cannot dictate in real time how those booking classes and ancillary services can be combined: if the airline wants to sell BAG1 and BAG2 to passengers booking in all three fare classes in general, the booking policy must always permit an optional BAG1 and BAG2 for classes Y, B, and M.[4] We term this constraint *fare class completeness*, and say that, to comply with traditional distribution system architecture, the

---

[4]It is possible that the airline marketing policies would offer some ancillaries complimentary to some fare classes, or would prohibit their purchase in other fare classes. These restrictions can be implemented through fare filing (and would affect the composition of M_k), but cannot be generated in real time.
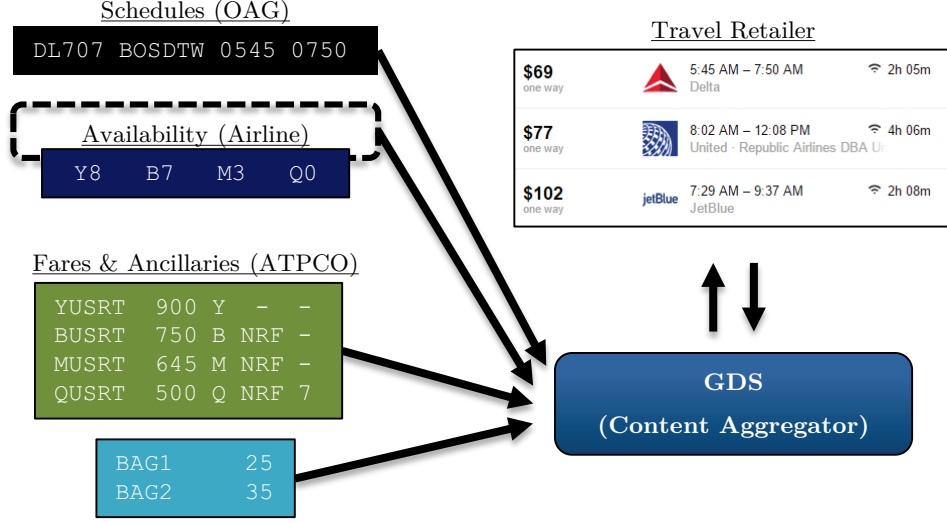
Figure 1: Schematic of a traditional distribution system, where consumers request travel options from a travel retailer, who then passes the request to a Global Distribution System. While the airline supplies schedules and fares/prices to OAG and ATPCO, the only real time request to the airline is for fare class availability. Search result image credit: https://www.google.com/flights.

offer sets considered by ACDP must be fare class complete.

Definition 1 (Fare Class Completeness). An offer set O is fare class complete (FCC) if and only if, for every fare class included in the set, all possible offers based on that fare class are also included in the set: if $(k, m) \in O$ for some $m \in M_k$, then O is FCC if and only if $(k, m') \in O$ for all $m' \in M_k$.

In addition, we impose the practical constraint that offer sets must be nested by fare order (a common assumption in many of today's revenue management systems):

Definition 2 (Nesting by Fare Order). An offer set O is nested by fare order (NFO) if and only if, for every fare class included in the set, all fare classes with a higher priced fare are also included in the set: if $(k, m) \in O$ for some $m \in M_k$, then O is NFO if and only if $(j, m') \in O$ for all $j \leq k$ and for some $m' \in M_j$.

Offer sets that are both fare class complete and nested by fare order are marketable by airlines within traditional distribution frameworks.

2.2. Ancillary Marginal Demand and Ancillary Marginal Revenue Heuristics

With the constraints described in the previous section, we can simplify our notation: the airline uses Equation 3 to choose a fare class complete, nested by fare order booking policy, which is equivalent to choosing a *lowest available fare class* with no RM-imposed limitations on ancillary purchase

options.[5] We denote this policy simply as k, which in our earlier notation is equivalent to the set $O = \{(i, m) \mid \forall i \leq k, \forall m \in M_i\}$. Following Fiig et al. (2010) and Walczak et al. (2010), we propose using AMD and AMR, which are optimal as inputs to an independent demand dynamic program, as heuristics to transform the demand and revenue inputs for existing (independent demand) *static* RM optimizers, such as EMSR. These systems are typically designed for fare-class level control, which is why we focus on the case where the airline must choose a fare class complete, nested by fare order policy to implement.

Using AMD and AMR as heuristic input modifiers provides an easy method to obtain the benefits of ancillary-awareness and choice-awareness without the need to significantly modify the core processes of the RM optimizer. The AMD demand corresponds to the incremental demand accommodated by opening class k, and the AMR fare to the incremental total revenue per unit of capacity earned by opening class k. We initially assume that each of these FCC/NFO policies is efficient, which allows the AMD and AMR outputs to give an adjusted demand and fare for each fare class, minimizing the need to modify the optimizer structure. We discuss the case of inefficient policies in Section 3.1. We express the total sale probability and total expected revenue for these booking policies as $TP_t(k)$ and $TR_t(k)$, and the heuristic versions of AMD and AMR are:

Ancillary Marginal Demand Transformation. The heuristic marginal demand $d_{k,t}$ associated with moving from fare class complete and nested by fare order booking policy $k - 1$ to k in time slice t is:

$$d_{k,t} = \lambda_t \left(TP_t(k) - TP_t(k-1) \right) \tag{8}$$

Ancillary Marginal Revenue Transformation. The heuristic marginal total revenue $f'_{k,t}$ associated with moving from fare class complete and nested by fare order booking policy $k - 1$ to k in time slice t is:

$$f'_{k,t} = \frac{TR_t(k) - TR_t(k-1)}{TP_t(k) - TP_t(k-1)} \tag{9}$$

In practice the airlines collect booking data, generate forecasts, and estimate parameters at various Data Collection Points (DCPs), which aggregate many time slices. In our simulations (Section 4), we break the booking window into 16 DCPs. We assume that choice probabilities and demand arrival rates are equal for each time slice within a DCP, and we express these parameters in terms of DCPs:

---

[5]To reiterate, ancillary purchase options for a given fare class could be limited by pre-specified marketing policies.

$$\lambda_t = \lambda_{dcp_t}, \quad P_{i,m,t}(k) = P_{i,m,dcp_t}(k), \quad TP_t(k) = TP_{dcp_t}(k)$$
$$TR_t(k) = TR_{dcp_t}(k), \quad f'_{kt} = f'_{k,dcp_t} \tag{10}$$

where dcp_t is the DCP that contains time slice t. We will typically refer to these quantities by DCP only, without reference to any particular time frame.

Gallego (2013) cautions that this heuristic approach is no longer an optimal solution, as the assumptions of the existing RM model are likely violated. For example, EMSR assumes that all low-fare demand arrives before high-fare demand, but under the general choice model we have incorporated there is no requirement that this assumption will hold. In addition, EMSR (typically) assumes normal demand distributions instead of Poisson demand distributions.

Despite the misalignment of these assumptions and loss of optimality, we believe (and our simulation results indicate) that the AMD and AMR heuristics can still provide a significant revenue benefit over traditional RM models when passengers make choices among fare classes and ancillary services.

2.3. Equivalence to Other Models

In this section we show that, under certain choice model conditions, the AMD and AMR heuristics are equivalent to the optimizer increment (OI) and OI combined with the (non-ancillary) marginal demand and revenue transformations.

The optimizer increment is a total revenue optimization heuristic for existing RM models based on supplying the optimizer with an adjusted fare the includes the expected ancillary revenue per passenger:

$$f'^{OI}_{k,dcp} = f_k + \bar{a}_{k,dcp} \tag{11}$$

where $f'^{OI}_{k,dcp}$ is the OI-adjusted fare for class k in DCP dcp and $\bar{a}_{k,dcp}$ is the expected ancillary revenue per booking in class k in DCP dcp. In practice, $\bar{a}$ would likely be estimated based on historical ancillary purchase data and could be aggregated across DCPs. No change is made to demand estimates.

The (non-ancillary) marginal demand and marginal revenue transformations (MD) and (MR) have the same structure as AMD and AMR Equations 8 and 9, except they rely on $TP^{MR}_{dcp}(k)$ and $TR^{MR}_{dcp}(k)$, which exclude the ancillary dimension:

$$TP^{MR}_{dcp}(k) = \sum_{i=1}^{k} \sum_{m \in M_k} P_{i,m,dcp}(k) \tag{12}$$

$$TR_{dcp}^{MR}(k) = \sum_{i=1}^{k} \sum_{m \in M_k} P_{i,m,dcp}(k) f_k \tag{13}$$

When these approaches are combined, the optimizer increment occurs before the expected revenue calculation:

$$TR_{dcp}^{OI+MR}(k) = \sum_{i=1}^{k} \sum_{m \in M_k} P_{i,m,dcp}(k) f_{k,dcp}^{\prime OI} = \sum_{i=1}^{k} \sum_{m \in M_k} P_{i,m,dcp}(k)(f_k + \bar{a}_{k,dcp}) \tag{14}$$

There is no additional change to total sale probability, so:

$$TP_{dcp}^{OI+MR}(k) = TP_{dcp}^{MR}(k) \tag{15}$$

2.3.1. Independent Demand Model

We first consider the independent demand choice model, in which each arriving customer has one preferred fare class k^* and combination of ancillary services m^*. If (k^*, m^*) is included in the airline's booking policy, the customer purchases it. Otherwise, they choose the no-fly option. We denote by $q_{k,m,dcp}$ the probability that an arriving consumer in DCP dcp has preference (k, m). Then, the independent demand model is specified by the choice probabilities:

$$P_{i,m,dcp}(k) = \begin{cases} q_{i,m,dcp} & i \le k \\ 0 & \text{otherwise} \end{cases} \tag{16}$$

Theorem 1. *With an independent demand model, the AMR heuristic has the same expected value as the optimizer increment: $f_{k,dcp}^{\prime AMR} = f_{k,dcp}^{\prime OI}$, where $f^{\prime AMR}$ is the AMR adjusted fare and $f^{\prime OI}$ is the optimizer increment adjusted fare.*

Proof. We first note that with the independent demand model $TP_{dcp}(k) = \sum_{i=1}^{k} \sum_{m \in M_i} q_{i,m,dcp}$ and $TR_{dcp}(k) = \sum_{i=1}^{k} \sum_{m \in M_i} q_{i,m,dcp}(f_i + a_{i,m})$. The expected ancillary revenue $\bar{a}_{k,dcp}$ from a booking in k class during DCP dcp is:

$$\bar{a}_{k,dcp} = \sum_{m \in M_k} a_{km} \Pr_{dcp}(\text{buy } m \mid \text{book } k)$$

$$= \sum_{m \in M_k} a_{km} \frac{\Pr_{dcp}(\text{buy } m \cap \text{book } k)}{\Pr_{dcp}(\text{book } k)}$$

$$= \sum_{m \in M_k} a_{km} \frac{q_{k,m,dcp}}{\sum_{m' \in M_k} q_{k,m',dcp}}$$

$$= \frac{\sum_{m \in M_k} a_{k,m} q_{k,m,dcp}}{\sum_{m \in M_k} q_{k,m,dcp}}$$

We can now show that the AMR fare is equal to the filed fare plus the (DCP-specific) expected ancillary revenue per passenger:

$$
\begin{aligned}
f'^{AMR}_{k,dcp} &= \frac{TR_{dcp}(k) - TR_{dcp}(k-1)}{TP_{dcp}(k) - TP_{dcp}(k-1)} \\
&= \frac{\sum_{i=1}^{k} \sum_{m \in M_i} q_{i,m,dcp}(f_i + a_{i,m}) - \sum_{i=1}^{k-1} \sum_{m \in M_i} q_{i,m,dcp}(f_i + a_{i,m})}{\sum_{i=1}^{k} \sum_{m \in M_i} q_{i,m,dcp} - \sum_{i=1}^{k-1} \sum_{m \in M_i} q_{i,m,dcp}} \\
&= \frac{f_k \sum_{m \in M_k} q_{k,m,dcp} + \sum_{m \in M_k} q_{k,m,dcp} a_{k,m}}{\sum_{m \in M_k} q_{k,m,dcp}} \\
&= f_k + \frac{\sum_{m \in M_k} a_{k,m} q_{k,m,dcp}}{\sum_{m \in M_k} q_{k,m,dcp}} \\
&= f_k + \bar{a}_{k,dcp}
\end{aligned}
\tag{17}
$$

Thus, $f'^{AMR}_{k,dcp} = f'^{OI}_{k,dcp}$ and the two formulations lead to equivalent adjusted fares. □

Remark. With the independent demand model, the ancillary marginal demand transformation has no effect on forecasting: the marginal demand associated with opening a class is the entire demand for the class, since passengers either purchase their (one) preferred class or do not fly.

2.3.2. Sequential Demand

We now consider the case when passengers exhibit sequential behavior, as defined by Bockelie and Belobaba (2017). Under the assumption that preferences about different types of fare class attributes and/or ancillary services are independent of each other, all sequential passengers booking in class i in DCP *dcp* have a constant conditional probability $w_{m|i,dcp}$ of purchasing ancillary combination m, regardless of which booking policy $k \geq i$ is in effect. Therefore, the choice probabilities

satisfy the following equation:

$$w_{m|i,dcp} = \text{Pr}_{dcp}(\text{buy } m \mid \text{book } i \text{ under policy } k)$$

$$= \frac{P_{i,m,dcp}(k)}{\sum_{m' \in M_i} P_{i,m',dcp}(k)} \quad \forall i \leq k \tag{18}$$

Theorem 2. *With a sequential demand model that satisfies Equation 18, the AMR heuristic has the same expected value as the optimizer increment combined with the (non-ancillary) Marginal Revenue transformation defined by Fiig et al. (2010) and Walczak et al. (2010): $f'^{AMR}_{k,dcp} = f'^{OI+MR}_{k,dcp}$, where f'^{AMR} is the AMR adjusted fare and f'^{OI+MR} is the optimizer increment with marginal revenue transformation adjusted fare.*

Proof. With sequential demand, the expected ancillary revenue in DCP dcp for a booking in class i is $\bar{a}_{i,dcp} = \sum_{m \in M_i} a_{i,m} w_{m|i,dcp}$; therefore $f'^{OI}_{k,dcp} = f_k + \sum_{m \in M_i} a_{i,m} w_{m|i,dcp}$. For the remainder of this section, we will refer to Equations 1 and 2 (and the DCP variants in Equation 10) as $TP^{AMR}_{dcp}(k)$ and $TR^{AMR}_{dcp}(k)$; the AMR adjusted fare is therefore:

$$f'^{AMR}_{k,dcp} = \frac{TR^{AMR}_{dcp}(k) - TR^{AMR}_{dcp}(k-1)}{TP^{AMR}_{dcp}(k) - TP^{AMR}_{dcp}(k-1)}$$

The optimizer increment/marginal revenue transformation adjusted fare has the same structure:

$$f'^{OI+MR}_{k,dcp} = \frac{TR^{OI+MR}_{dcp}(k) - TR^{OI+MR}_{dcp}(k-1)}{TP^{OI+MR}_{dcp}(k) - TP^{OI+MR}_{dcp}(k-1)}$$

To prove $f'^{AMR}_{k,dcp} = f'^{OI+MR}_{k,dcp}$, we show that the formulations have equivalent total sale probabilities and total expected revenues: $TP^{AMR}_{dcp}(j) = TP^{OI+MR}_{dcp}(j)$ and $TR^{AMR}_{dcp}(j) = TR^{OI+MR}_{dcp}(j)$. Total sale probability for the marginal revenue transformation is given by Equation 15:

$$TP^{OI+MR}_{dcp}(j) = \sum_{i=1}^{j} P_{i,dcp}(j) = \sum_{i=1}^{j} \sum_{m \in M_i} P_{i,m,dcp}(j)$$

which is equivalent to $TP^{AMR}_{dcp}(j)$ as defined in Equations 1 and 10. Next we show that the total expected revenues are equivalent:

$$
\begin{aligned}
TR^{OI+MR}_{dcp}(k) &= \sum_{i=1}^{k} P_{i,dcp}(k) f'^{OI}_{i,dcp} \\
&= \sum_{i=1}^{k} \left(\left(\sum_{m \in M_i} P_{i,m,dcp}(k) \right) \left(f_i + \sum_{m \in M_i} a_{i,m} w_{m|i,dcp} \right) \right) \\
&= \sum_{i=1}^{k} \left(f_i \left(\sum_{m \in M_i} P_{i,m,dcp}(k) \right) + \frac{\sum_{m \in M_i} a_{i,m} P_{i,m,dcp}(k)}{\sum_{m' \in M_i} P_{i,m',dcp}(k)} \sum_{m \in M_i} P_{i,m,dcp}(k) \right) \\
&= \sum_{i=1}^{k} \sum_{m \in M_i} P_{i,m,dcp}(k)(f_i + a_{im}) \\
&= TR^{AMR}_{dcp}(k)
\end{aligned}
$$

Since $TP^{AMR} = TP^{OI+MR}$ and $TR^{AMR} = TR^{OI+MR}$ for the sequential choice model, the two approaches have the same expected adjusted fares (and same adjusted demands). □

Under a general choice model, Equation 18 will not hold: the probability that a passenger purchases a particular ancillary combination, even given that they book in class k, may vary based on the other classes offered. In this case, AMR and the optimizer increment plus marginal revenue transformation will lead to different adjusted fares and, potentially, different booking policies. Our AMD and AMR formulations, by explicitly including the ancillary dimension of passenger choice, can more precisely measure the marginal total revenue (or demand) associated with opening a fare class than the previous approaches.

3. Operationalization

In this section we discuss additional processes and assumptions that are necessary or helpful to implement the AMD and AMR heuristics. First, as noted above, the AMD and AMR transformations described in Equations 8 and 9 are only valid if the booking policies k and $k-1$ are optimal solutions to the original Equation 3. As described by Talluri and van Ryzin (2004), the solutions to Equation 3 must always be *efficient sets*, and it is possible that the booking policies k and/or $k-1$ are not efficient. Borrowing from previous Fare Family research, we propose several convex hull approximation methods, which we will refer to as "gap-filling" methods, to cope with inefficient booking policies (see Hopperstad (2008) and Fiig et al. (2012)). Second, the ACDP formulation assumes that demand arrival rates λ_{dcp} are known; in reality, forecasting demand is a significant challenge for airlines. We develop an AMD Forecasting Model below to help cope with this challenge, particularly when the heuristics are coupled with an optimizer that requires a forecast of demand-to-come by fare class.

3.1. Gap-Filling

In cases where a fare class complete, nested by fare order policy k is not efficient, AMR adjusted fares will be inverted (i.e. the adjusted fare for class k may be *less than* the adjusted fare for class $k + 1$).

The optimal approach to dealing with these inefficient policies is, of course, to not offer them. However, our proposed use case for the AMD and AMR heuristics is to feed adjusted demands and adjusted fares to an existing RM optimizer (such as EMSR) to provide ancillary and choice-awareness without significantly revising the core of the optimization procedure. As existing RM optimizers are based on fare class demands and revenues, it is important that the output of AMD/AMR can be expressed in terms of fare classes as well: each class needs an AMD demand and an AMR fare.

We propose three different mechanisms, which we term *gap-filling* methods, to deal with inefficient booking policies while maintaining compatibility with the structure of existing RM optimizers. The first two, vertical and horizontal gap-filling, involve approximating the airline's computed choice probabilities to move the inefficient policies onto the efficient frontier. The third mechanism, exclusion gap-filling, involves strategically modifying the AMD/AMR outputs to prevent the optimizer from producing an inefficient booking policy.

The outputs of gap-filling are the adjusted total sale probabilities $TP'_{dcp}(k)$ and adjusted total expected revenues $TR'_{dcp}(k)$ for each DCP and for each booking policy. These adjusted TP' and TR' values are used in subsequent AMD/AMR processes.

To illustrate each of these mechanisms, we consider a running example with six fare classes and one ancillary service (corresponding to DCP 2 of the simulation results presented in Section 4). The computed total sale probabilities and total expected revenues are listed in the upper-left portion of Table 1, along with the associated EMSR booking limits. Note that without any gap-filling, the AMR adjusted fares for FC 2 and FC 5 are inverted.

Figure 2 shows a plot of the efficient frontier and the fare class complete/nested by fare order booking policies $\{1\}$, $\{1, 2\}$, and $\{1, 2, 3\}$; the policy of offering classes 1 and 2 is inefficient. Graphically, the AMR adjusted fare is the slope of the line segment connecting two adjacent policies. Vertical and horizontal gap-filling operate by shifting the policy FC 1–2 until it falls on the segment connecting FC 1 and FC 1–3. Vertical gap-filling shifts the policy up; horizontal gap-filling shifts the policy left. Other (diagonal) gap-filling policies would be possible, but are not investigated here.

3.1.1. Vertical Gap-Filling

With the vertical gap-filling approximation, the airline's computed total expected revenue $TR_{dcp}(k)$ is increased for inefficient policies until the policy falls on the efficient frontier. AMD forecasting is
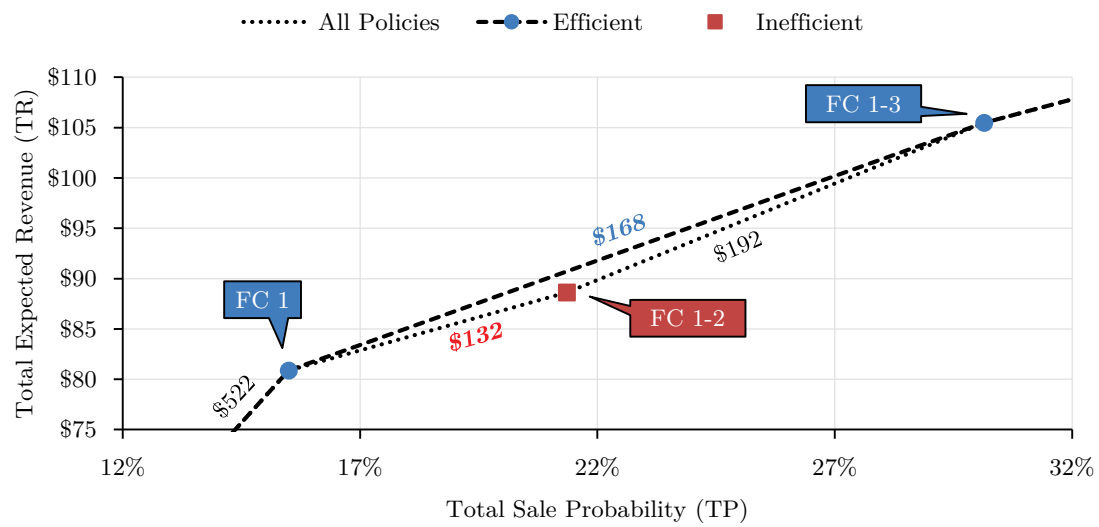
16

Figure 2: Portion of an example convex hull, showing efficient policies in blue circles and inefficient policies in red squares. The slope (AMR adjusted fare) of various segments is indicated, with emphasis in red for inverted fares and in blue for gap-filled fares.

Table 1: Example AMD mean demand $\tilde{\mu}_k$, AMR fares f'_k, and associated EMSR booking limits (BL) for DCP 2 with various gap-filling methods. Inverted fares are emphasized in red; changes from no gap-filling emphasized in blue with the size of the change indicated in parentheses. (Forecast volume mean $\tilde{\mu} = 85$, forecast volume standard variance $\tilde{\sigma}^2 = 900$, and capacity remaining $x = 110$).

	No Gap-Filling					Vertical Gap-Filling				
k	$TP(k)$	$TR(k)$	$\tilde{\mu}_k$	f'_k	BL	$TP'(k)$	$TR'(k)$	$\tilde{\mu}_k$	f'_k	BL
1	15.5%	$81	13	$522	110	15.5%	$81	13	$522	110
2	21.4%	$89	5	*$132*	102	21.4%	*$91 (+2.3%)*	5	*$168 (+27%)*	102
3	30.1%	$105	7	$192	101	30.1%	$105	7	*$168 (-12%)*	*100 (-1)*
4	52.8%	$134	19	$126	96	52.8%	$134	19	$126	96
5	71.7%	$148	16	*$77*	85	71.7%	*$149 (+0.4%)*	16	*$80 (+4%)*	*86 (+1)*
6	100.0%	$172	24	$82	78	100.0%	$172	24	*$80 (-2%)*	78

	Horizontal Gap-Filling					Exclusion Gap-Filling				
k	$TP(k)$	$TR(k)$	$\tilde{\mu}_k$	f'_k	BL	$TP'(k)$	$TR'(k)$	$\tilde{\mu}_k$	f'_k	BL
1	15.5%	$81	13	$522	110	15.5%	$81	13	$522	110
2	*20.1% (-1.2pt)*	$89	*4 (-21%)*	*$168 (+27%)*	102	*n/a*	*n/a*	*0 (-100%)*	*$168 (+27%)*	102
3	30.1%	$105	*9 (+14%)*	*$168 (-12%)*	101	30.1%	$105	*12 (+67%)*	*$168 (-12%)*	*102 (+1)*
4	52.8%	$134	19	$126	96	52.8%	$134	19	$126	96
5	*71.0% (-0.7pt)*	$148	*15 (-4%)*	*$80 (+4%)*	*86 (+1)*	*n/a*	*n/a*	*0 (-100%)*	*$80 (+4%)*	85
6	100.0%	$172	*25 (+2%)*	*$80 (-2%)*	78	100.0%	$172	*40 (+67%)*	*$80 (-2%)*	*85 (+7)*

not directly affected by this process, as the total sale probabilities remain unchanged. It is important to note that this approximation changes only the airline's perception of expected revenue; the airline does not change fares, fare class restrictions, or ancillary prices, so actual passenger choice will not be affected.

The vertical gap-filling approximations are:

$$TP'_{dcp}(k) = TP_{dcp}(k)$$

$$TR'_{dcp}(k) = \begin{cases} TR_{dcp}(k) & \text{if } k \text{ is efficient} \\ TR_{dcp}(k-1) + \frac{TR_{dcp}(k+1)-TR_{dcp}(k-1)}{TP_{dcp}(k+1)-TP_{dcp}(k-1)} \left(TP_{dcp}(k) - TP_{dcp}(k-1)\right) & \text{otherwise} \end{cases}$$

The result of vertical gap-filling is shown in upper-right section of Table 1. With the example parameters, vertical gap-filling increases TR for FC 2 by 2.3% and for FC 5 by 0.4%. These small changes, however, translate into a 27% increase in the adjusted fare of FC 2, as well as smaller changes for the adjusted fares of FC 3, 5, and 6. In addition, the EMSR booking limits change; one fewer seat is permitted to FC 3 and one additional seat is permitted to FC 5.

3.1.2. Horizontal Gap-Filling

With the horizontal gap-filling approximation, the total sale probability $TP_{dcp}(k)$ for inefficient policies is decreased until the policy lies on the efficient frontier. AMD forecasting is directly affected by this process. While the airline does not adjust $TR_{dcp}(k)$, because of the change to TP_{dcp}, AMR adjusted fares also change. The AMR fares from horizontal gap-filling will be equal to those of vertical gap-filling because adjusted fares are dictated by the properties of the efficient booking policies in both cases.

The horizontal gap-filling approximations are:

$$TP'_{dcp}(k) = \begin{cases} TP_{dcp}(k) & \text{if } k \text{ is efficient} \\ TP_{dcp}(k-1) + \frac{TP_{dcp}(k+1)-TP_{dcp}(k-1)}{TR_{dcp}(k+1)-TR_{dcp}(k-1)} \left(TR_{dcp}(k) - TR_{dcp}(k-1)\right) & \text{otherwise} \end{cases}$$

$$TR'_{dcp}(k) = TR_{dcp}(k)$$

19

In our running example, shown in the bottom-left of Table 1, the change to total sale probabilities due to horizontal gap-filling is small—a decrease of 1.2 percentage points for FC 2, and of 0.7 points for FC 5. However, these small changes again have large impacts on AMD and AMR: adjusted demand decreases 21% for FC 2 and increases 24% for FC 3; there are also smaller changes for FC 5 and FC 6. The net result of horizontal gap-filling, in this example, is an increase by 1 in the FC 5 booking limit compared to the no gap-filling case. Recall that vertical gap-filling increased the FC 5 booking limit by 1, but also reduced the FC 2 booking limit.

3.1.3. Exclusion Gap-Filling

Exclusion gap-filling is the most mathematically-correct approach to dealing with inefficient policies when AMD and AMR are used as heuristic input modifiers for existing RM optimizers. Exclusion gap-filling is a multistage process that attempts to produce AMD demands and AMR fares that prevent the RM optimizer from selecting inefficient booking policies.

In the first stage, inefficient policies are discarded, and AMD demands and AMR fares are computed using only the efficient policies. Next, demands and adjusted fares are filled in for the inefficient policies: a demand of zero and an adjusted fare equal to the adjusted fare of the next efficient policy. This process ensures that, at the time of optimization with EMSR, the inefficient policies (in the case of this example, offering FC 1–2 or FC 1–5) will never be selected.

It is important to note that "FC 1–2" as an inefficient policy does not mean that consumers should be prohibited from booking in FC 2; when FC 1–3 are available (which is an efficient policy, as shown in Figure 2), consumers are free to choose to buy-up to FC 2 if they wish.

The bottom-right of Table 1 shows the effect of exclusion gap-filling on our running example. Note that the AMR fares are equal to those produced by vertical and horizontal gap-filling. However, the AMD demands differ from both methods; the exclusion gap-filling forecast has less high-fare demand and more lower-fare demand (note that all demand from FC 2 in the no gap-filling case gets moved to FC 3 with exclusion gap-filling; the same applies for FC 5 and 6). This leads exclusion gap-filling to have less aggressive booking limits than the no, vertical, and horizontal gap-filling: compared to no gap-filling, exclusion gap-filling increases the FC 6 booking limit by 7, and increases the FC 3 booking limit by 1.

The adjusted total sale probability and adjusted total expected revenue from any of the gap-filling methods will be used to compute AMR fares and AMD demands.

3.2. Demand Forecasting

Our AMD Forecasting Model provides an approach for estimating parameters for the distribution of demand volume, based on historical bookings. This model is an extension of Q-Forecasting

([Hopperstad and Belobaba, 2004](#)) to support generic fare structures and generic passenger choice models. We propose a four step process, in which we convert historical booking observations into estimates of historical demand volume, and then forecast future demand based on the historical volume estimates. Note that this process occurs after any gap-filling.

1. Convert observed (historical) bookings to equivalent "Q"-bookings, which represents the number of bookings that would have been received in the past if all fare classes and all ancillary services had been available. We assume that the airline has recorded which offer set was presented to the consumer for each booking. The equivalent Q-bookings for DCP dcp on previous departure day dep is given by:

$$qb_{dcp,dep} = \sum_{k=1}^{n_{FC}} \frac{b_{k,dcp,dep}}{TP'_{dcp}(k)} \tag{19}$$

where n_{FC} is the number of fare classes, $b_{k,dcp,dep}$ is the number of bookings received in DCP dcp on previous departure date dep when FCC and NFO booking policy k was offered to consumers, and $qb_{dcp,dep}$ is the equivalent Q-bookings for DCP dcp on previous departure date dep, and serves as the estimate for total demand volume for that DCP and day.

2. Detruncate historical observations for any instances where all classes were closed:

$$\hat{qb}_{dcp,dep} = d(qb_{dcp,dep}, Z) \tag{20}$$

where $\hat{qb}_{dcp,dep}$ is the detruncated (unconstrained) equivalent Q-bookings for DCP dcp on previous departure date dep, $d()$ is a detruncation function, and Z is a vector of other data or parameters for the detruncation process, including whether or not all classes were closed. For details on detruncation methods, see [Lee (1990)](#), [Wickham (1995)](#), [Skwarek (1996)](#), [Weatherford and Pölt (2002)](#), and [Queenan et al. (2007)](#).

3. Forecast future demand volume distribution parameters based on detruncated historical equivalent Q-demand for n_{dep} previous departure days:

$$\mu_{dcp} = \frac{1}{n_{dep}} \sum_{dep=1}^{n_{dep}} \hat{qb}_{dcp,dep}$$

$$\sigma_{dcp}^2 = \frac{1}{n_{dep}-1} \sum_{dep=1}^{n_{dep}} (\hat{qb}_{dcp,dep} - \mu_{dcp})^2$$

where μ_{dcp} is the forecast future demand volume mean for DCP dcp and σ_{dcp}^2 is the forecast future demand volume variance for DCP dcp .

4. Partition demand within each DCP to each booking policy:

$$\mu_{k,dcp} = \mu_{dcp}(TP'_{dcp}(k) - TP'_{dcp}(k-1))$$

$$\sigma_{k,dcp}^2 = \sigma_{dcp}^2(TP'_{dcp}(k) - TP'_{dcp}(k-1))$$
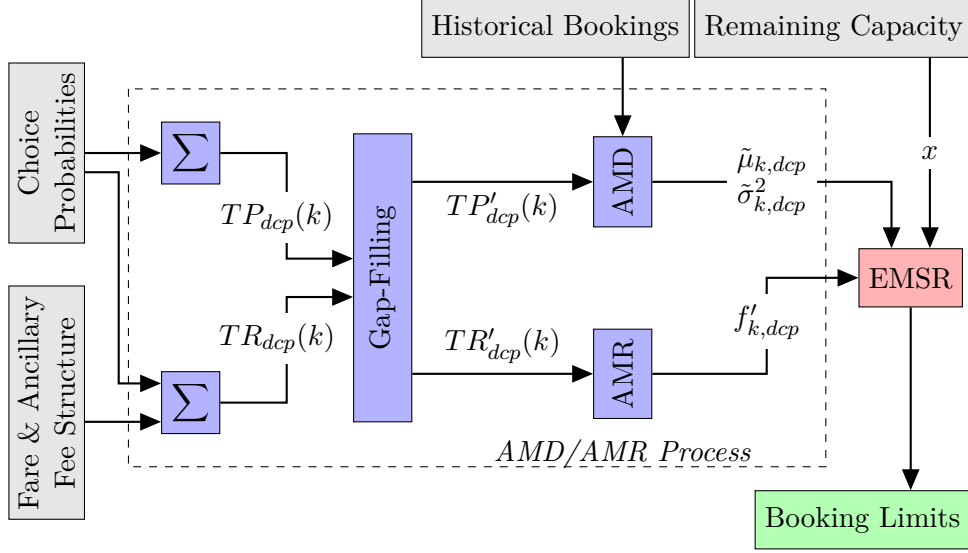
21

Figure 3: AMD/AMR process diagram when used as a heuristic in conjunction EMSR. Inputs are shown in grey, AMD/AMR processes and operations in blue, the RM optimizer in red, and the RM output in green.

For RM optimizers such as EMSR that utilize demand-to-come forecasts, instead of DCP-specific forecasts, the demands must be aggregated across DCPs. The forecast of all future demand to come, generated at the start of DCP dcp, for policy k, is:

$$\tilde{\mu}_{k,dcp} = \sum_{i=dcp}^{n_{DCP}} \mu_{k,i} \qquad \tilde{\sigma}_{k,dcp}^2 = \sum_{i=dcp}^{n_{DCP}} \sigma_{k,i}^2$$

and the total future demand to come from the start of DCP dcp (aggregating across all booking policies) is:

$$\tilde{\mu}_{dcp} = \sum_{k=1}^{n_{FC}} \tilde{\mu}_{k,dcp} \qquad \tilde{\sigma}_{dcp}^2 = \sum_{k=1}^{n_{FC}} \tilde{\sigma}_{k,dcp}^2$$

We use these aggregated-across-DCP demand parameters ($\tilde{\mu}_{k,dcp}$ and $\tilde{\sigma}_{k,dcp}^2$) in our implementation of AMD and AMR with EMSR.

3.3. Process Summary

An overview of the complete AMD and AMR methodology, when used as a heuristic with EMSR, is illustrated in Figure 3. The process requires input choice probabilities ($P_{i,m,dcp}(k)$) for each of the fare class complete/nested by fare order booking policies, fares and ancillary prices, and historical bookings. The process returns AMR fares and AMD demand estimates (mean and variance) for each class, generated at the start of each DCP. These adjusted fares and demands are fed to the RM optimizer (along with remaining flight capacity), which then returns a booking policy in the form of booking limits.

4. Simulated Performance

In this section we utilize the Passenger Origin-Destination Simulator (PODS) to assess the performance of our AMD and AMR heuristics. We provide summary results for a range of parameter values, and detailed results for one illustrative case. In all cases we compare the performance of AMD/AMR against previous ideas for total revenue optimization.

4.1. The Passenger Origin-Destination Simulator

The Passenger Origin-Destination Simulator is a numerical simulator designed to model the interactions between passengers and airline revenue management systems, and is used to evaluate the performance of various RM forecasting and optimization models. PODS consists of two interdependent modules, as shown in Figure 4. The Passenger Module generates consumers within origin-destination (OD) markets and assigns randomly drawn preferences and budgetary constraints to each passenger. Consumers evaluate the itineraries, fare classes, and ancillary services available to them at the time they are generated and, based on their preferences, book the option that is most appealing (or do not fly if no options are within their budgetary constraint). Bockelie and Belobaba (2017) describe the PODS consumer choice models in detail.

The Airline Module collects data on historical bookings, uses the data to forecast demand for future flights, and then runs a revenue management optimizer to set booking policies for the network. A key aspect of PODS is that airlines do not have access to the underlying demand generation parameters; they must forecast demand based on observed bookings from previous (simulated) departures. Each airline in the simulation maintains its own booking database, forecasting, and optimization systems. The decisions made by passengers affect the historical booking databases and thus forecasting, optimization, and future availability decisions by the airlines. The decisions made by airlines affect the availability given to passengers, and thus the future booking decisions.

We will use PODS to study the effects of the AMD and AMR heuristics for total revenue management. Recent examples of other PODS studies include assessing the performance and impacts of dynamic pricing (Wittman and Belobaba, 2018) and the influence of RM system users (Weatherford, 2016).

4.2. Single Airline, Single Flight Network

We utilize a single airline, single flight leg network within PODS for our studies. The flight has a capacity of 130 seats. The airline offers one optional ancillary service and six economy fare classes, denoted FC 1 (the most expensive and least restricted, $500) through FC 6 (the least expensive, most restricted, and subject to advance purchase requirements, $125). The airline divides its
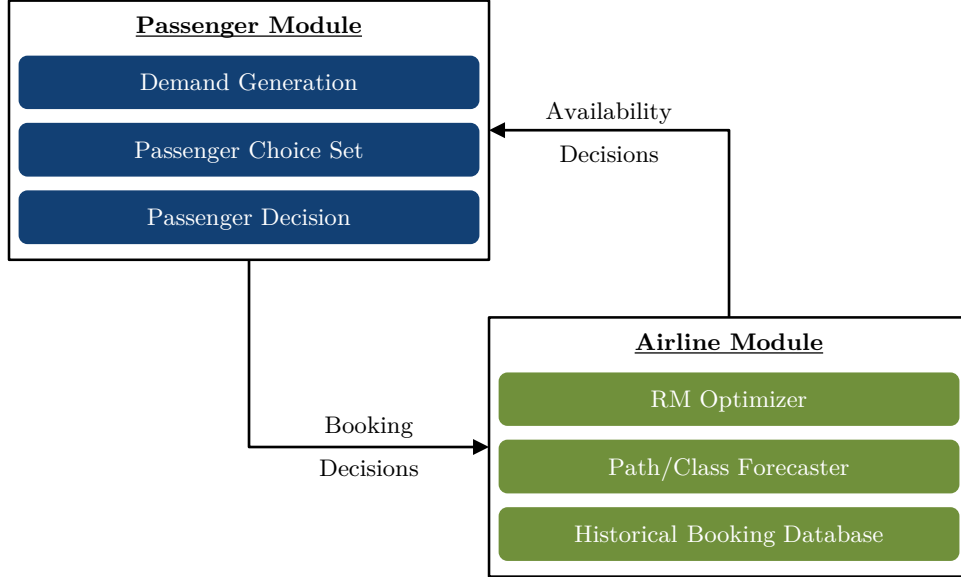
Figure 4: Schematic of the PODS Simulator. Adopted from M. Wittman.

booking window into 16 DCPs, with the end of DCP 16 corresponding to departure. Forecasts are generated and booking limits are re-optimized at the start of each DCP.

This network features two consumer demand segments, business and leisure. Business passengers tend, although do not always, to have higher budgetary constraints, to book closer to departure, and to be more averse to fare class restrictions (such as a Saturday-night stay, or non-refundability). The average business passenger and leisure passenger booking curves, as well as the business/leisure mix, are shown in Figure 5. Note that the early DCPs have a low proportion of business travelers shopping, while shoppers in the later DCPs are predominantly from the business segment. Because each PODS simulation covers many departure days, small changes in revenues (about 0.1%) or bookings (about 0.1 per class) are statistically significant.

4.3. Experimental Outline

In all of our simulations, the airline uses EMSRb as its RM optimizer. As a baseline, we consider the case where the airline optimizes based on filed fares and uses an independent-demand forecasting model (see Belobaba and Weatherford (1996) and Littlewood (1972)). For AMD and AMR, we assume that the airline has (accurately) estimated a passenger choice model, and can therefore compute the probabilities $P_{i,m,dcp}(k)$ as described in Equation 10. Unless otherwise noted, the airline will employ Exclusion Gap-Filling when using AMD and AMR (because it shows the highest revenue in our simulations, as shown in Section 4.4.1).

We will compare the performance of AMD and AMR against the baseline as well as three existing approaches for accounting for ancillary revenue and/or passenger choice: the optimizer increment
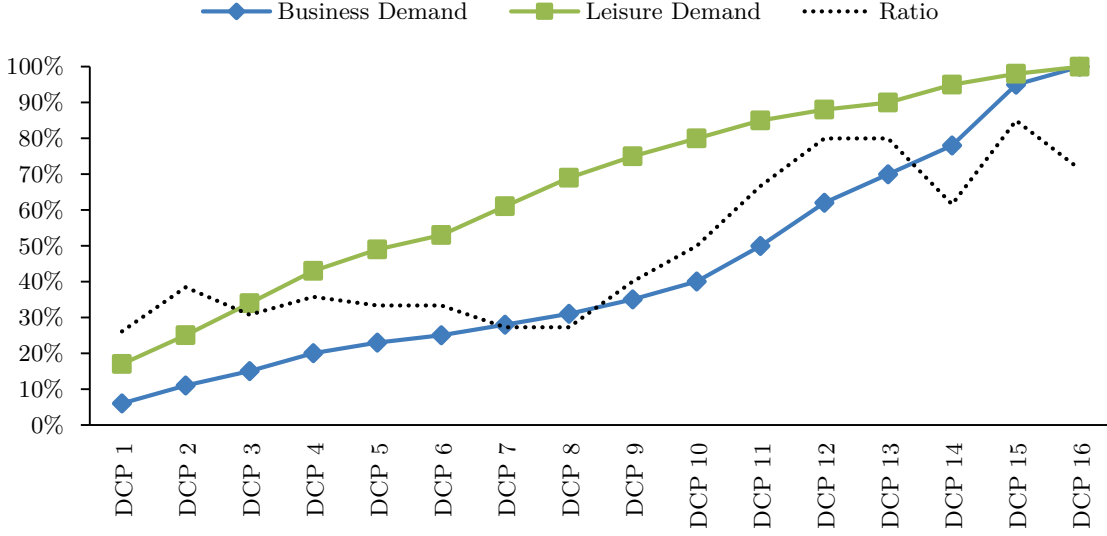
24

Figure 5: Cumulative average business and leisure consumer arrival curves, and the ratio of average business to leisure arrivals within each DCP.

(OI), hybrid forecasting and fare adjustment (HF/FA), and a combination of the two approaches (OI + HF/FA). In our implementation of OI, the airline estimates $\bar{a}_{k,dcp}$ based on historical purchases aggregated across all DCPs. Recall that hybrid forecasting and fare adjustment are operationalized versions of the (non-ancillary) marginal revenue and marginal demand transformations, where demand is divided into two groups: product-oriented, which is forecasted with an independent demand model and no marginal revenue transformation, and price-oriented, which is forecasted with a marginal demand model and has a marginal revenue transformation applied to the fares (based on a negative exponential sell-up curve); the final demand and fare values sent to the optimizer are a combination of the price and product values.

We focus our detailed assessment of AMD/AMR on a representative case in which the ancillary service has a price of \$50, and both consumer segments have a mean valuation for the service of \$50. We assume that passengers in the simulation are aware of ancillary prices and incorporate ancillary preferences into their decision making process ("simultaneous" behavior, as defined in Bockelie and Belobaba (2017)).

The total probabilities $TP_{dcp}(k)$ computed by the airline are shown in Figure 6. Early in the booking window, the airline calculates that the probability of sell-up is low (as indicated by the low total sale probability for FC 1, and reflecting the low proportion of business travelers). However, later in the booking window, when the portion of business passengers is higher, the total sale probability for higher-value classes increases, suggesting that more sell-up is possible. Note that the total sale probability for FC 6 is always 100%, reflecting that all generated passengers in the simulation can afford to purchase the lowest published fare.
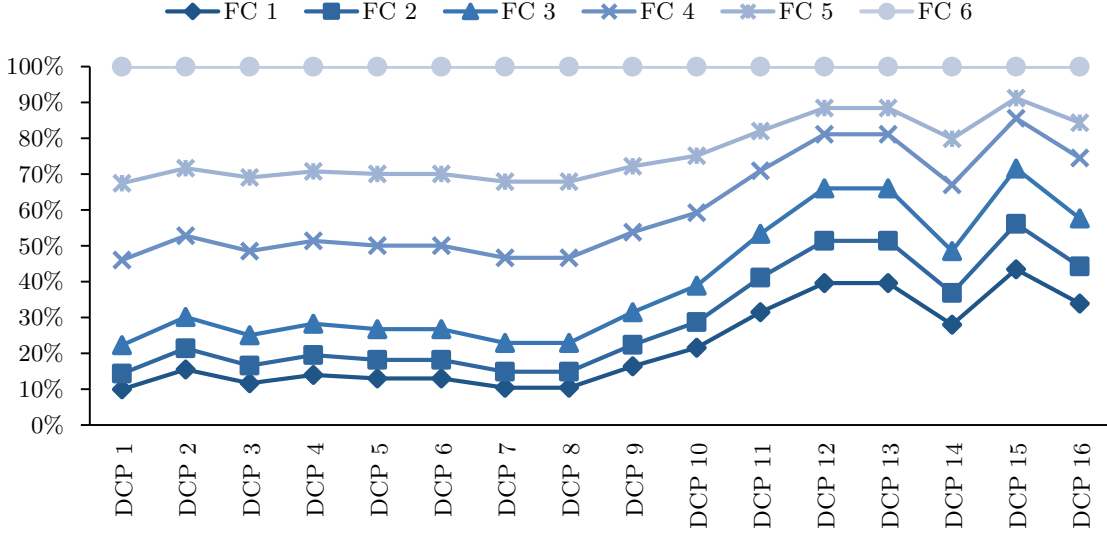
Figure 6: Total sale probability $TP_{dcp}(k)$ for each class as computed by the airline (prior to any gap-filling).

4.4. Results

Table 2 lists booking and ancillary purchase data for the baseline simulation. In the baseline scenario, about 33% of passengers purchase the ancillary service, with a much higher purchase rate in the higher value fare classes (46% for FC 1) than in the lower value fare classes (25% in FC 6); likewise, the average ancillary revenue per booking is highest in FC 1 ($23) and lowest in FC 6 ($12). Despite the lower ancillary purchase rate and lower average ancillary revenue, because fares are lower in FC 6 than FC 1, the portion of total revenue derived from ancillary sales is highest in the lower value classes (11% in FC 5) and lowest in higher value classes (4% in FC 1). For reference, major US airlines report about 8% of total revenue from ancillary services, according to the US Department of Transportation.[6]

The lower ancillary purchase rate in the lower value classes is driven by a fundamental behavioral assumption in the Simultaneous choice model (Bockelie and Belobaba, 2017): passengers have an overall budgetary constraint that limits their spending on the combination of fare and ancillary services. Passengers booking in the lower value (and highly restricted) fare classes tend to have lower budgets, which constrains their ability to afford ancillary services. These basic ancillary purchase and revenue trends in the baseline case are similar in the other experimental cases.

The total revenue, load factor, and yield (expressed as revenue per passenger mile) for each of the experimental cases (with AMD/AMR using exclusion gap-filling) are shown in Table 3. AMD/AMR produces a revenue increase of 1.8% over baseline; HF/FA produces a gain of 1.2%, 0.6 pts less

---

[6]US DOT Form 41, Schedule P-1.2

Table 2: Baseline booking and ancillary purchase data.

	FC 1	FC 2	FC 3	FC 4	FC 5	FC 6	Total
Fare	$500	$390	$295	$200	$160	$125	$203
Average ancillary per passenger	$23	$22	$21	$22	$19	$12	$17
Portion of total revenue from ancillary	4%	5%	7%	10%	11%	9%	8%
Ancillary purchase rate	46%	44%	41%	43%	39%	25%	33%
Bookings	6	14	10	8	16	56	109

Table 3: Revenue, load factor, and yield for the baseline case and experimental cases.

	Baseline	OI	HF/FA	OI + HF/FA	AMD/AMR
Total Revenue	$23,906	$23,876	$24,181	$24,171	$24,336
Load Factor	83.8%	83.9%	82.5%	82.6%	82.5%
Total Yield	21.95	21.90	22.54	22.50	22.70
Change from baseline					
Total Revenue		-0.1%	+1.2%	+1.1%	+1.8%
Load Factor		+0.1 pts	-1.3 pts	-1.2 pts	-1.3 pts
Total Yield		-0.2%	+2.7%	+2.5%	+3.4%

than AMD/AMR. Both AMD/AMR and HF/FA decrease load factor by 1.3 pts, and both increase total yield, although the increase is larger for AMD/AMR (3.4%) than for HF/FA (2.7%).

The optimizer increment has a small negative effect on revenue: an 0.1 pt decrease compared to baseline, and a reduction of the benefit of HF/FA by 0.1 pt. OI increases load factor by 0.1 pt when used alone, and decreases the load factor loss due to HF/FA by 0.1 pt. Although the revenue and load factor changes due to OI are small, they are directionally consistent with the results seen in numerous studies within the MIT PODS Research Consortium (e.g. Bockelie and Belobaba (2016)).

These revenue and load factor changes are driven by shifts in the booking mix, as shown in Figure 7. Both HF/FA and AMD/AMR reduce bookings in the lowest value class (FC 6), while increasing bookings in the highest value class (FC 1). The methods differ in the magnitudes of changes: AMD/AMR reduces bookings in FC 6 by about 2, while HF/FA reduces by 5. Changes in bookings in higher value fare classes have a disproportionate effect on total revenue: recall that FC 1 has a fare of $500, and FC 6 has a fare of $125, so each FC 1 booking is worth four times the ticket revenue of an FC 6 booking. The FC 1 booking increase with AMD/AMR is worth six times the revenue loss associated with the FC 6 booking decrease.

While both AMD/AMR and HF/FA reduce FC 6 bookings (which saves space for later arriving, higher value FC 1 bookings), only AMD/AMR also reduces FC 5 bookings; HF/FA leads to an increase in FC 5 (as well as FC 4). The booking changes by fare class due to OI are minimal, and the changes with OI + HF/FA are approximately equal to the sum of the changes in the OI case
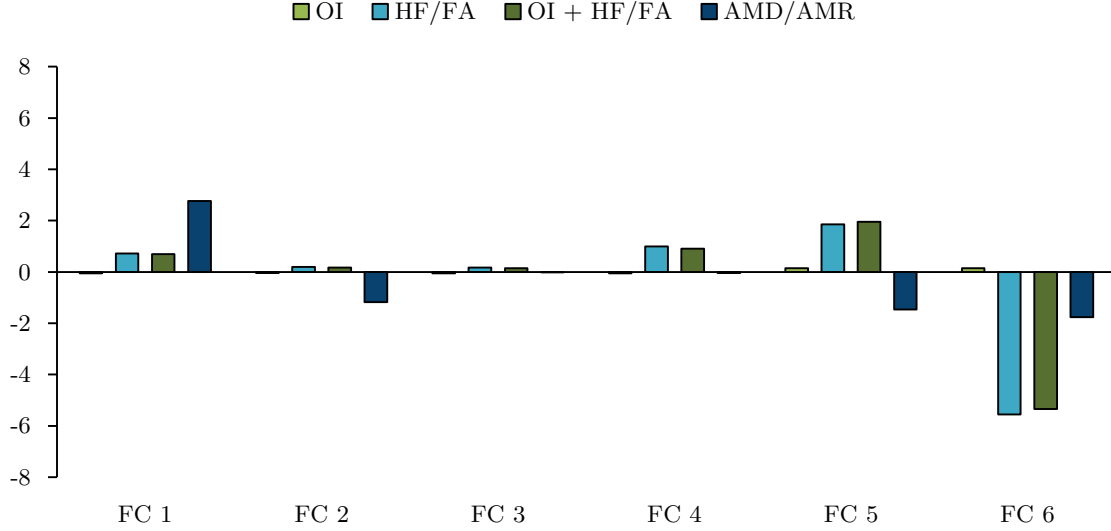
Figure 7: Change in bookings vs baseline by fare class due to OI, HF/FA, OI + HF/FA, and AMD/AMR.

and in the HF/FA case.

The initial demand and ticket revenue forecast (generated at the start of DCP 1) is shown in Figure 8. Compared to the baseline, AMD provides a lower demand forecast mean (by 15%), but a higher demand forecast standard deviation (by 33%). Overall, the probability that demand exceeds the aircraft capacity of 130 (a rough indicator of whether the capacity constraint should restrict availability) is 24% for AMD, a reduction from the 39% of the baseline or 34% of HF/FA. Despite the lower volume of demand with AMD, however, the value of demand with AMD is greater because the composition of the AMD forecast is shifted toward FC 1. AMD produces an initial forecast of ticket revenue (computed as a sum of demand for each class multiplied by the fare for each class) 13% higher than the baseline, with a standard deviation 65% higher. A higher value forecast, for the same demand mean and standard deviation and same optimizer fares, will lead to more aggressive availability decisions and fewer low-value bookings.

In addition to the higher value forecast, the AMR adjusted fares are lower than filed fares, especially for the low value fare classes. The higher value forecast of AMD combined with the reduced optimizer fares of AMR reduces availability of lower-value fare classes, as shown in Figure 9. Reducing lower class availability forces consumers to buy-up to higher value classes, reducing load factor and increasing yields (and in this case increasing total revenue). By explicitly accounting for ancillary revenue and passenger choices, AMD/AMR can more precisely close classes. Note that AMD/AMR has a smaller reduction in FC 6 availability than HF/FA, but a larger reduction in FC 5 availability (especially in DCP 7 and 8), leading to the booking shifts seen in Figure 7.
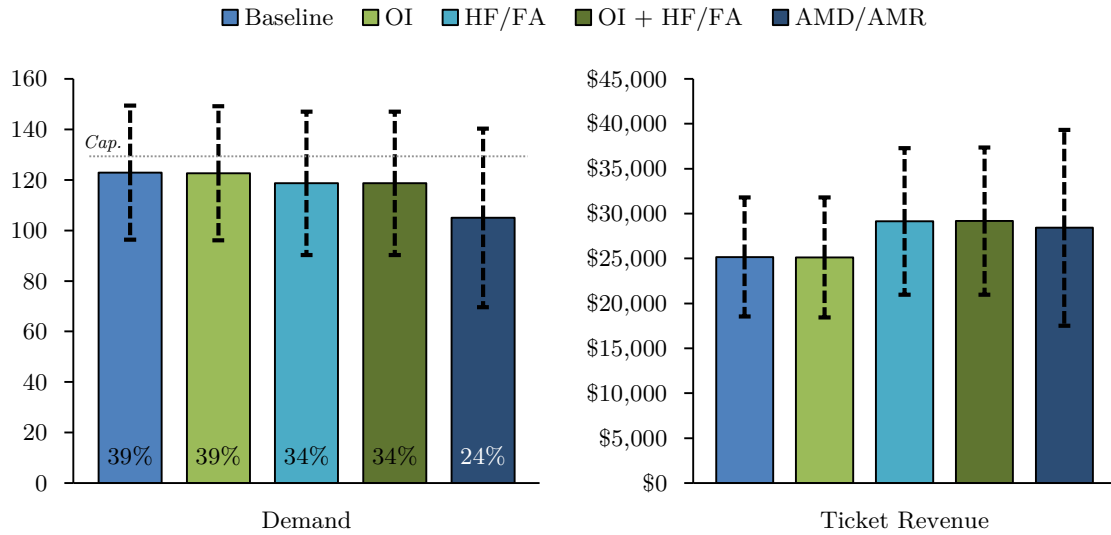
Figure 8: Initial total demand and ticket revenue forecasts (generated at the start of DCP 1). Solid bars show forecast mean; dashed error lines show forecast standard deviation. The forecast probability that demand exceeds capacity (of 130) is indicated at the base of the solid bars.
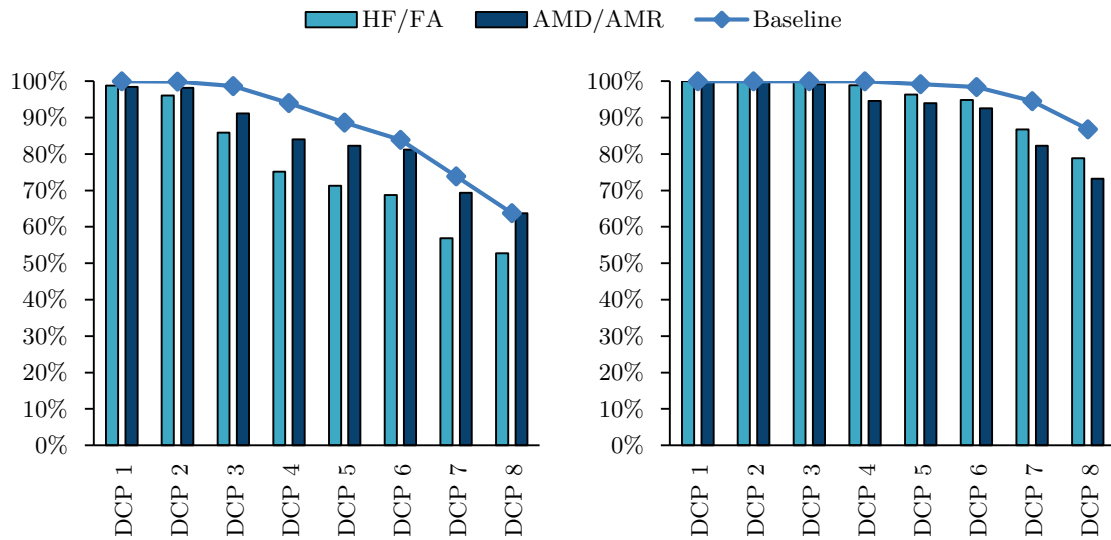

Figure 9: Availability (measured as the portion of time a class is available for sale) for FC 6 (left) and FC 5 (right) for first 8 DCPs for the baseline, HF/FA, and AMD/AMR cases.

Table 4: Revenue, load factor, and yield for the baseline case and each AMD/AMR gap-filling case.

	Baseline	None	Vertical	Horizontal	Exclusion
Total Revenue	$23,906	$24,248	$24,299	$24,321	$24,336
Load Factor	83.8%	80.7%	81.5%	81.9%	82.5%
Total Yield	21.95	23.12	22.92	22.85	22.70
Change from baseline					
Total Revenue		+1.4%	+1.6%	+1.7%	+1.8%
Load Factor		-3.1 pts	-2.2 pts	-1.9 pts	-1.3 pts
Total Yield		+5.3%	+4.4%	+4.1%	+3.4%

4.4.1. Effect of Gap-Filling

Table 4 lists performance data for each of the gap-filling mechanisms. Exclusion gap-filling produces a revenue increase of 1.8% over baseline; AMD/AMR with the other gap-filling methods has a smaller revenue gain, with the worst revenue performance when gap-filling is not used. Load factor changes with AMD/AMR are inverse to total revenue changes: the largest revenue gain (exclusion gap-filling) has the smallest load factor loss (-1.3 pts), while the lowest revenue gain (1.4%, no gap-filling) has the largest load factor loss (-3.1 pts). All of the gap-filling methods have higher revenue than HF/FA.

As shown in Figure 10, exclusion gap-filling, which has the highest revenue and is most mathematically correct, has the smallest reduction on FC 6 bookings; the other gap-filling methods all have FC 6 booking reductions of similar magnitude to HF/FA. In addition, exclusion gap-filling is the only method that reduces FC 5 bookings; the other methods trade large FC 6 losses for smaller FC 5 gains. This is an expected result; as shown in Table 1, exclusion gap-filling protects no seats for FC 5, and therefore FC 5 and FC 6 have the same booking limit. In that example, exclusion gap-filling has an FC 6 booking limit 7 seats greater than any of the other gap-filling methods, include no gap-filling. Despite the greater booking limit for FC 6, though, exclusion gap-filling has the second greatest increase in FC 1 bookings vs baseline (the additional space to accommodate FC 1 customers is provided by accepting fewer bookings in FC 5 and 4 compared to the other gap-filling mechanisms).

The differences in fare class booking changes amongst the gap-filling methods illustrate that no gap-filling is the most aggressive form of AMD/AMR, followed by vertical gap-filling, then horizontal gap-filling, and finally exclusion gap-filling. This variation in aggressiveness is a function of the AMD forecast generated by each approach: recall that vertical, horizontal, and exclusion gap-filling all have the same AMR fares. Horizontal gap-filling partially reduces demand forecasts for inefficient policies (and increases the forecast for the next efficient policy); exclusion gap-filling completely eliminates inefficient policy forecasts, and shifts all demand to the next efficient policy. Thus, exclusion gap-filling will always have less aggressive availability, and will accept more FC 6
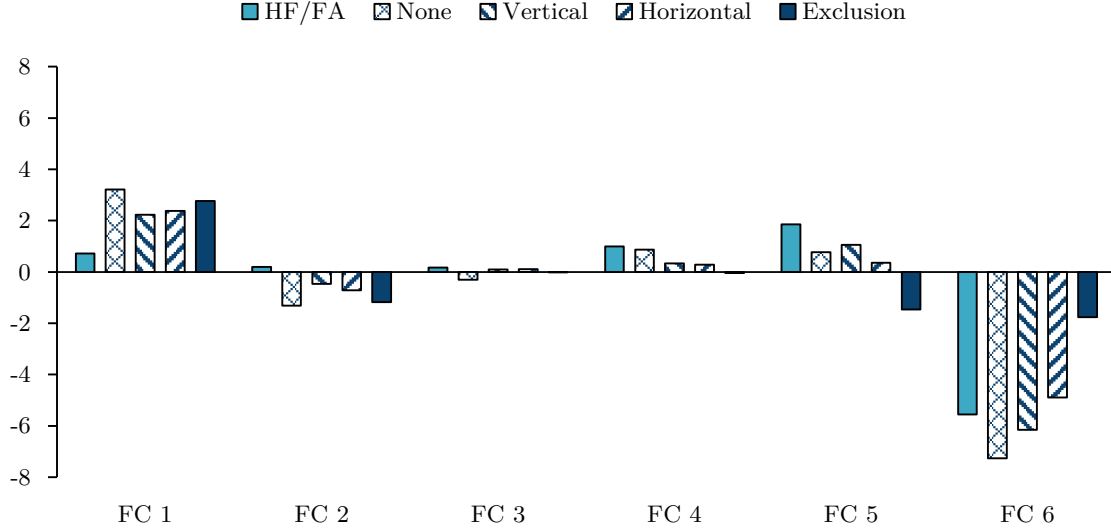
Figure 10: Change in bookings by fare class vs baseline due to HF/FA and AMD/AMR with various gap-filling settings.

bookings than the other gap-filling methods. The increase in FC 6 availability (relative to other gap-filling) means less space is protected for FC 5 bookings, which produces the decrease in FC 5 (relative to both the baseline and other gap-filling methods) seen in Figure 10.

4.4.2. Sensitivity Assessment

In this section we assess the sensitivity of our results to the input ancillary prices and ancillary utility parameters. We vary ancillary prices from $25 to $100, with mean passenger ancillary utilities equal to 75%, 100%, or 125% of the ancillary price. We consider three combinations of utilities: a leisure-oriented service, where leisure passenger mean utility is 125% of the ancillary price and business passenger utility is 75% of the ancillary price, a business-oriented service where the percentages are reversed, and an equally appealing service, with a mean utility of 100% of the ancillary price for both segments.

The results of the previous section are consistent across the range of parameters tested. The change in total revenue (vs baseline) is shown in Figure 11. The optimizer increment leads to revenue losses on the order of 0.1% in all cases. Hybrid forecasting and fare adjustment provides revenue increases of 1.0–1.2% over baseline in each of the cases. AMD and AMR provide an additional 0.6–0.7 pts of revenue benefit over HF/FA, for a total gain of about 1.8% over baseline when the ancillary service is optional in all classes. The benefit of AMD and AMR over HF/FA or baseline is relatively stable over the range of ancillary prices and ancillary utilities tested. Figure 12 lists the airline's load factor for each simulation; in general, the optimizer increment slightly increases load factor, while HF/FA and AMD/AMR decrease load factor compared to the baseline. HF/FA has a slightly

Figure 11: Percent change in total revenue vs baseline. The horizontal axis lists ancillary prices (\$25–\$100) and utility configurations.

greater load factor decrease (around -1.2 pts) than AMD/AMR (around -0.9 pts) vs baseline.

Ancillary purchase rates, average ancillary revenue by passenger, and the portion of total revenue generated by ancillary services (by class and overall) vary widely across these parameter ranges. However, the effect of AMD and AMR on booking mix is similar across price ranges, as shown in Figure 13, but has more variability as relative utilities change (from the leisure-oriented to equally appealing to business-oriented cases). The leisure-oriented utilities result in a decrease in bookings in all but the highest fare class. The business-oriented utilities decrease bookings in FC 2, 5 and 6, but increase bookings in FC 3 and 4; the magnitude of the FC 6 decrease (and the total decrease across all fare classes) is greater with the business-oriented utilities. The business-oriented utilities have a larger decrease in FC 6 for two reasons. First, with the higher business utility, more passengers booking in later TFs and more passengers booking in higher value fare classes will purchase the ancillary service, making it more important to save space for late arriving, high value customers. Second, with the lower leisure utility, fewer passengers booking in early TFs and fewer passengers booking in lower value classes will purchase the ancillary, making the value of early FC 6 bookings low, and compounding the booking limit effects of the higher business utility.

5. New Distribution Capability

Industry efforts are underway to reduce, or eliminate, the fare class completeness constraint described in Section 2.1. The International Air Transport Association (IATA) is leading the development of New Distribution Capability (NDC), which is a suite of new distribution technologies and standards. One relevant aspect of NDC is that it would enable airlines to display and sell more content
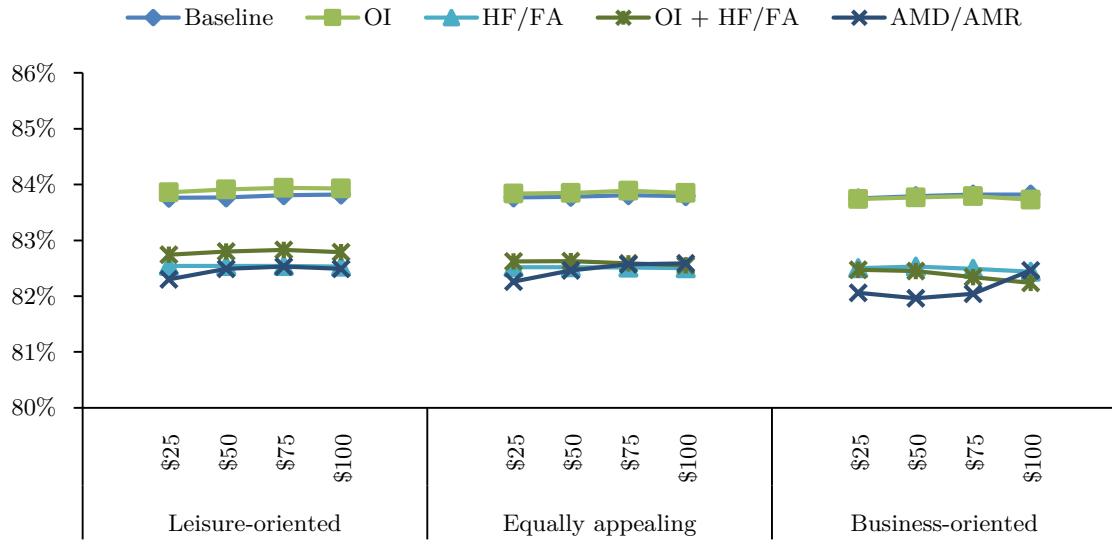
Figure 12: Load factor. The horizontal axis lists ancillary prices ($25–$100) and utility configurations.
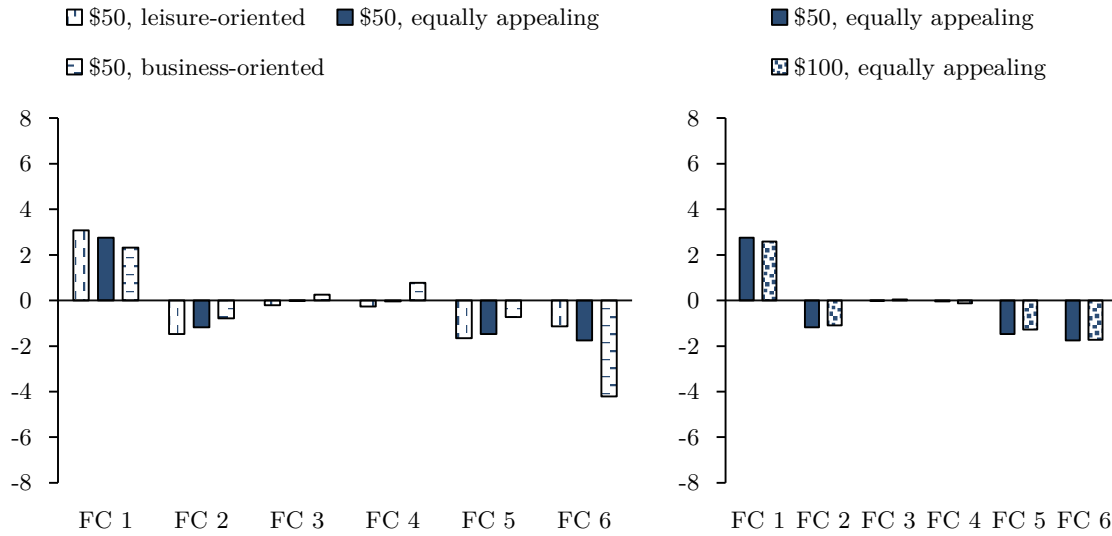

Figure 13: Change in bookings by fare class vs relevant baseline due to AMD/AMR with exclusion gap-filling. The left chart shows various utility settings and a $50 ancillary price; the right chart shows various ancillary prices when both consumer segments have a mean ancillary utility equal to price. The $50 equally appealing case corresponds to the results in the previous portion of this section.

through more channels, particularly indirect channels. When NDC is implemented by an airline and its distribution partners, GDSs would no longer be required to aggregate schedules, availability, and fares to assemble sets of booking options, as described in Figure 1. Instead, GDSs (or other content aggregators) could request offers directly from airlines.[7] Most importantly for this work, these offers will no longer be limited to fare class availability. The airline could respond to each request with a specifically designed offer or set of offers—each offer would comprise of an itinerary, a set of zero or more ancillary services, various purchase/use restrictions, and a price.

The shift from a traditional distribution environment to NDC has significant implications for total revenue optimization. Because NDC increases the ability of airlines to display and sell ancillary services in indirect channels, airlines expect ancillary revenues to rise with the implementation of NDC. In addition, because NDC moves away from fare class-centered availability control, airlines will have to develop offer generation systems. Finally, NDC allows more detailed booking requests (such as indicating round trip travel, or frequent flyer status), which could enable offer generation systems to create *individualized* sets of offers. At the limit, each offer could be personally and dynamically constructed and priced for each consumer.

We imagine, at least in the near term, a more limited view of NDC that allows airlines to control the availability of specific fare class and ancillary service combinations in real time, but still relies on filed fares and filed ancillary prices and does not personalize or individualize offers. In such an environment, the Ancillary Choice Dynamic Program we have developed in this paper could serve as the basis for an offer generation system. In such an implementation, the solutions to Equation 3 need not be fare class complete or potentially even nested by fare order. The solution space for the problem would grow significantly, and the task of determining which offer sets are efficient would be more complex. However, significant gains in total revenue could also be possible. As one example, consider a case where demand is very high relative to capacity. With a traditional distribution system, the airline's highest expected unit revenue booking policy is to offer only FC 1. However, with NDC and offer generation, it would be possible to create an even higher unit revenue offer: FC 1 with a requirement that the passenger purchase any typically optional ancillary services. While requiring passengers to purchase ancillary services (likely marketed as a bundle, with a total price equal to the fare plus the ancillary fees) would reduce the probability of a consumer buying the offer, in very high demand to capacity scenarios (i.e. when the bid price is very high), it is a revenue-maximizing strategy. Additional development and experimentation is required to assess the revenue potential of this approach.

---

[7]Airlines and GDSs could still choose to have the GDS assemble offers.

6. Conclusions and Future Work

In this paper, we have developed a new dynamic programming model for total revenue optimization that incorporates fares, ancillary revenues, and passenger choices. The model produces an optimal set of offers, which we define as a fare class and a combination of ancillary services, to be presented at any given time. Following previous work on choice-based RM, we use the Ancillary Marginal Demand transformation and Ancillary Marginal Revenue transformation to convert our ancillary and choice-aware DP into an equivalent independent demand formulation. After addressing practical distribution constraints, we devised a series of processes to utilize AMD and AMR as total revenue optimization heuristics in conjunction with existing RM optimizers, and we developed the Ancillary Marginal Demand forecasting model to provide demand volume estimates.

We then simulated the performance of our heuristics in a single airline, single flight leg environment with a range of ancillary prices and passenger ancillary utilities. We found that AMD and AMR, when used with EMSR, can produce revenue gains of approximately 1.8%, where previous choice-aware methods such as hybrid forecasting and fare adjustment only increased revenue by 1.2%, and the previous total revenue optimization heuristic (optimizer increment) decreased total revenue. The gains from AMD and AMR were driven by a reduction of availability in lower-value classes, increasing space for later-arriving, higher-value customers, and forcing other customers to buy-up to higher-value fare classes.

Numerous extensions and enhancements to our work are possible. For example, while exact network formulations for dynamic programs suffer from exploding dimensionality, our heuristics could be extended to support a network setting. A key challenge will be determining the level of detail and specificity necessary in the input choice probabilities to maintain reasonable revenue performance; the extent to which these probabilities could be aggregated and/or scaled across different markets and fare structures is unknown. In addition, competitive effects in a network setting raise questions about the degree to which competitor offerings should be explicitly incorporated in the model: failing to account for competitor offerings could lead to availability decisions that are too aggressive, while explicitly modeling competitor actions increases data and computation requirements.

Our work has assumed that the airline has knowledge of conditional choice probabilities; in reality, these would need to be estimated. Further work is required to develop efficient estimation methods, and to understand how inaccuracies in the input choice probabilities affect revenue performance.

Finally, as addressed in Section 5, the rise of New Distribution Capability could allow airlines to have significantly more control over the offers they produce. We believe that this work could be extended to generate offers based on filed fares and prices. A separate, potentially larger challenge, would be devising *dynamic* offer generation engines that also incorporate a dynamic pricing aspect.

Acknowledgments

Ancarani, F., Gerstner, E., Posselt, T., Radic, D., 2009. Could higher fees lead to lower prices? Journal of Product & Brand Management 18 (4), 297–305.

Balcombe, K., Fraser, I., Harris, L., 2009. Consumer willingness to pay for in-flight service and comfort levels: a choice experiment. Journal of Air Transport Management 15, 221–226.

Belobaba, P., Weatherford, L., 1996. Comparing decision rules that incorporate customer diversion in perishable asset revenue management situations. Decision Science 27, 343–363.

Bockelie, A., Belobaba, P., 2016. Incorporating ancillary services in passenger choice: Revenue and booking impacts.
URL https://agifors.wildapricot.org/resources/Documents/RM/2016/program.html

Bockelie, A., Belobaba, P., 2017. Incorporating ancillary services in airline passenger choice models. Journal of Revenue and Pricing Management 16 (6), 553–568.

Boyd, A., Kallesen, R., 2004. The science of revenue management when passengers purchase the lowest available fare. Journal of Revenue and Pricing Management 3 (2).

Bratu, S., 1998. Network value concept in airline revenue management. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

Brueckner, J. K., Lee, D. N., Picard, P. M., Singer, E., 2015. Product unbundling in the travel industry: The economics of airline bag fees. Journal of Economics and Management Strategy 24, 457–484.

Espino, R., Martín, J. C., Román, C., 2008. Analyzing the effect of preference heterogeneity on willingness to pay for improving service quality in an airline choice context. Transportation Research Part E 44, 593–606.

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

Fiig, T., Isler, K., Hopperstad, C., Skytte Olsen, S., 2012. Forecasting and optimization of fare families. Journal of Revenue and Pricing Management 11 (3), 322–342.

Gabaix, X., Laibson, D., 2006. Shrouded attributes, consumer myopia, and information suppression in competitive markets. Quarterly Journal of Economics 121 (2), 505–540.

Gallego, G., 2013. Single resource revenue management problems with dependent demands. Working Paper.

Hao, E., 2014. Ancillary revenues in the airline industry: Impacts on revenue management and distribution systems. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

Hopperstad, C., 2008. Q forecasting with fare families – methodology.

Hopperstad, C., Belobaba, P., 2004. Alternative RM algorithms for unrestricted fare structures. In: Proceedings of the 2004 AGIFORS RM Study Group Meeting.

Lautenbacher, C., Stidham, S., 1999. The underlying Markov decision process in the single-leg airline yield management problem. Transportation Science 33, 136–146.

Lee, A. O., 1990. Airline reservations forecasting: Probabilistic and statistical models of the booking process. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

Littlewood, K., 1972. Forecasting and control of passenger bookings. In: Proceedings of the 12th AGIFORS Symposium.

Metters, R., Queenan, C., Ferguson, M., Harrison, L., Higbie, J., Ward, S., Barfield, B., Farley, T., Kuyumcu, H. A., Duggasani, A., May–Jun 2008. The "killer application" of revenue management: Harrah's Cherokee Casino & Hotel. Interfaces 38 (3), 161–175.

Mumbower, S., Garrow, L. A., Newman, J. P., 2015. Investigating airline customers' premium coach seat purchases and implications for optimal pricing strategies. Transportation Research Part A 73, 53–69.

Phillips, R., 2005. Pricing and Revenue Optimization. Stanford Business Books, Stanford, CA.

Queenan, C. C., Ferguson, M., Higbie, J., Kapoor, R., 2007. A comparison of unconstraining methods to improve revenue management systems. Production and Operations Management 16 (6), 729–746.

Scotti, D., Dresner, M., 2015. The impact of baggage fees on passenger demand on US air routes. Transport Policy 43, 4–10.

Shulman, J. D., Geng, X., 2013. Add-on pricing by asymmetric firms. Management Science 59 (4), 899–917.

Skwarek, D. K., 1996. Competitive impacts of yield management system components: Forecasting and sell-up models. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

Smith, B., Penn, C., 1988. Analysis of alternate origin-destination control strategies. In: Proceedings of the 28th AGIFORS Symposium.

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

Taubmann, H., 2016. One year since Lufthansa's distribution cost charge—what we've learned.
URL http://www.amadeus.com/blog/29/06/one-year-since-lufthansas-distribution-cost-charge-weve-learned/

Walczak, D., Mardan, S., Kallesen, R., 2010. Customer choice, fare adjustments and the marginal expected revenue data transformation: A note on using old yield management techniques in the brave new world of pricing. Journal of Revenue and Pricing Management 9 (1-2), 94–109.

Weatherford, L., 2016. Tests of revenue management performance under different demand correlation assumptions. Journal of Revenue and Pricing Management 15 (5), 399–416.

Weatherford, L., Pölt, S., 2002. Better unconstraining of airline demand data in revenue management systems for improved forecast accuracy and greater revenues. Journal of Revenue and Pricing Management 1 (3), 234–254.

Wickham, R., 1995. Evaluation of forecasting techniques for short-term demand of air transportation. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

Wittman, M. D., Belobaba, P. P., 2018. Customized dynamic pricing of airline fare products. Journal of Revenue and Pricing Management 17 (2), 78–90.

Zhuang, W., Li, M. Z., 2012. Dynamic pricing with two revenue streams. Operations Research Letters 40 (1), 46–51.

Zou, L., Yu, C., Rhoades, D., Waguespack, B., 2017. The pricing responses of non-bag fee airlines to the use of bag fees in the US air travel market. Journal of Air Transport Management 65, 209–219.