Premium cabin capacity sharing strategies: airline RM insights

Premium cabin capacity sharing strategies: airline RM insights

Cite this article as: Germán Escovar-Álvarez and Peter P. Belobaba, Premium cabin capacity sharing strategies: airline RM insights, Journal of Revenue and Pricing Management https://doi.org/10.1057/s41272-021-00326-x

Premium Cabin Capacity Sharing Strategies: Airline RM Insights

Germán Escovar-Álvarez[a,1] and Peter P. Belobaba*[a,2]

* Corresponding author

[a] Massachusetts Institute of Technology, International Center for Air Transportation,
77 Massachusetts Avenue, Building 33-318, Cambridge, MA 02139
Email: gescovar@alum.mit.edu; belobaba@mit.edu

[1] holds a Master of Science in Transportation from MIT, where his research focused on revenue management strategies for airlines operating dual cabin aircraft. He has worked as a consultant on transportation, supporting airlines, airports, investors, and public agencies worldwide.

[2] is Principal Research Scientist at MIT, where he teaches graduate courses on the Airline Industry and Airline Management. He is Program Manager of MIT's Global Airline Industry Program and Director of the MIT PODS Revenue Management Research Consortium. He holds a Master of Science degree and a PhD in Flight Transportation Systems from MIT. He has worked as a consultant on revenue management systems at over 50 airlines and other companies worldwide.

Premium Cabin Capacity Sharing Strategies: Airline RM Insights

Germán Escovar-Álvarez and Peter P. Belobaba

Abstract

This paper explores revenue management strategies for premium cabin capacity sharing, that is, accepting additional bookings in economy fare classes when some premium cabin seats are expected to remain empty. Two heuristics are presented and tested to evaluate the impacts of premium cabin capacity sharing on airline revenues and loads. These RM strategies are tested using the Passenger Origin Destination Simulation (PODS), which replicates the interactions between airline revenue management systems and passenger choice of airline itineraries, fare types and service offerings.

Based on the simulations, the methodologies proposed can generate total revenue gains of up to 1.1%. However, two caveats are identified: first, losses in the revenue captured from premium fare classes are likely to be experienced due to displacement by economy fare class passengers. Second, premium cabin capacity sharing should only be implemented in the final stages of the booking process; otherwise, the sharing heuristics can result in revenue losses for the airline.

Keywords: airline revenue management, cabin capacity sharing, premium versus economy cabin booking policy, cabin optimization

Abstract

1. Introduction

A carrier using two cabins in its aircraft has the opportunity to share the capacity of its premium cabin, allowing for passengers that purchase economy fare class tickets to occupy seats in the premium cabin when these are not expected to be filled by passengers paying premium cabin fares.

This paper explores the impacts that cabin-sharing revenue management strategies can have on the revenue performance of an airline. Our focus is on the simulated impacts on both total and economy versus premium cabin loads and revenues, given the use of realistic revenue management (RM) systems, and passengers that choose between the economy and premium cabin options offered by the airline. The following two sections review, respectively, the relevant literature and the premium cabin capacity sharing mechanisms considered in this paper. Then, we briefly describe the PODS simulator used in this study. Section five presents the results of the simulations, considering leg-based and OD control RM systems. We conclude with a discussion of the implications of the simulation results for the implementation of premium cabin capacity sharing mechanisms.

2. Literature Review

The multiple cabin revenue management problem aims to maximize the total revenue of an aircraft with multiple cabins by sharing the seats of a higher service cabin with passengers who book in the fare classes corresponding to a cabin that offers lower service standards. The main assumption for this approach is that passengers booked in the economy cabin will accept to be upgraded to a cabin that offers better service.

Alstrup et al. (1986) develop one of the earliest models found in the literature that splits the passengers into two types, “euro-class” and “tourist-class”. Passengers can be either upgraded from tourist to euro class or downgraded from euro to tourist class (the airline assumes a cost if a passenger is downgraded). The result of their model is a booking policy that defines the maximum number of reservations that can be accepted at a certain time of the booking period for each kind of passenger, given the number of passengers already booked in each cabin.

Lepage (2013) focuses on the single-leg multiple cabin revenue management problem; his thesis explores six different algorithms developed and evaluated using the Passenger Origin-Destination Simulator ("PODS"). Two of the evaluated strategies are based on a dynamic programming (DP) formulation and four are based on heuristics. Based on the simulations, Lepage found that revenue increases resulting from these approaches can be as high as 2.5% (compared to distinct cabin EMSR leg control), but gains were observed in all the tested scenarios. Overall, the heuristics outperformed the DP methods.

Walczak (2010) hypothesizes that greater flexibility in the mechanisms for sharing the premium cabin capacity results in higher revenue for the airlines. For the problem in which the capacity of each cabin is fixed, two DP formulations are proposed. The performance of the DP formulations is compared with five heuristics and it is found that, among these, the ones that perform better are the ones that use a joint bid price vector to control availability once the economy cabin is sold out.

Gallego and Stefanescu (2009) highlight the role that upgrades can play when there is a mismatch between supply and demand, as in the cases in which there is large premium cabin capacity in a market with low business passenger demand. One of the main results presented is that optimality could be lost if fairness in the upgrade process was to be guaranteed (e.g. if only highest-fare economy class was allowed to be accommodated in premium cabin) because the seller losses flexibility to maximize its revenues.

Finally, Rauch and Poelt (2015) proposed a Simple Capacity Sharing Heuristic for Fixed Compartments. It considers distinct bid prices vectors for both economy and premium cabins, and an additional bid price vector that considers jointly premium and economy demand and capacity (as if there was just a single compartment). The heuristic proposes controlling demand for premium fare classes by the effective premium bid price, defined as the maximum of the premium bid price and the joint bid price. On the other hand, demand for economy fare classes is controlled by the effective economy bid price, defined as the minimum of the economy bid price and the joint bid price (when economy cabin is full, the effective economy bid price is the joint bid price).

This study explores the multiple cabin revenue management problem focusing on the case of airlines operating dual cabin aircraft. Instead of splitting passengers into categories as in Alstrup et al. (1986), passengers in our simulations make their decisions based on their characteristics (e.g.

willingness to pay, disutilities) and on the options made available by the airlines. This paper considers the leg-based Full EMSR heuristic proposed by Lepage (2013) and proposes time-frame protection mechanisms to overcome its weaknesses. Furthermore, it builds on the methodology proposed by Rauch and Poelt (2015) to evaluate the performance of a premium cabin capacity sharing mechanism with OD-control RM. More detailed descriptions of both the algorithms and simulation results can be found in the MIT Master's thesis of Escovar Alvarez (2016).

3. Premium Cabin Capacity Sharing Mechanisms

The premium capacity sharing methodologies explored in this paper assume that passengers who book in economy fare classes would be pleased to enjoy the higher service quality offered in the premium cabin. However, we need to make two clarifications: First, economy cabin seats are not shared with passengers that book in premium fare classes since these passengers would not accept downgrades. Second, the upgrades discussed in this paper refer only to the effort of an airline to increase revenue by having additional flexibility to allocate passengers to aircraft seats; it does not consider other kinds of upgrades, such as those associated with frequent flyer programs, and sales of paid upgrades, etc.

3.1.1 Capacity Sharing with Leg-Control RM

As our baseline for comparisons, we apply leg-control EMSRb distinctly to each of the cabins, as if these were different aircraft.

The proposed leg-control premium cabin capacity sharing approach is based on the Shared Nesting "Full EMSR" methodology described by Lepage (2013); it consists of applying EMSRb to the entire aircraft capacity including both cabins and modifying the booking limits for premium fare classes based on the capacities of each cabin. This version of Full EMSR includes premium fare class demands in the calculation of lower nested fare classes in both cabins. This suggests a logical inconsistency in the heuristic, since premium fare class bookings cannot be downgraded to the economy cabins under our assumptions. Lepage (2013) also tested a variant of Full EMSR in which only economy fare class demands are used to set economy cabin booking limits, but found its performance to suffer even more from opening up lower fare classes in the economy cabin. We therefore explore the Full EMSR algorithm as originally described.

Figure 1 illustrates this mechanism by identifying that the number of premium cabin seats shared is equivalent to the difference between the booking limit protection level for the highest economy fare class (in our case FC5) and the premium cabin capacity.

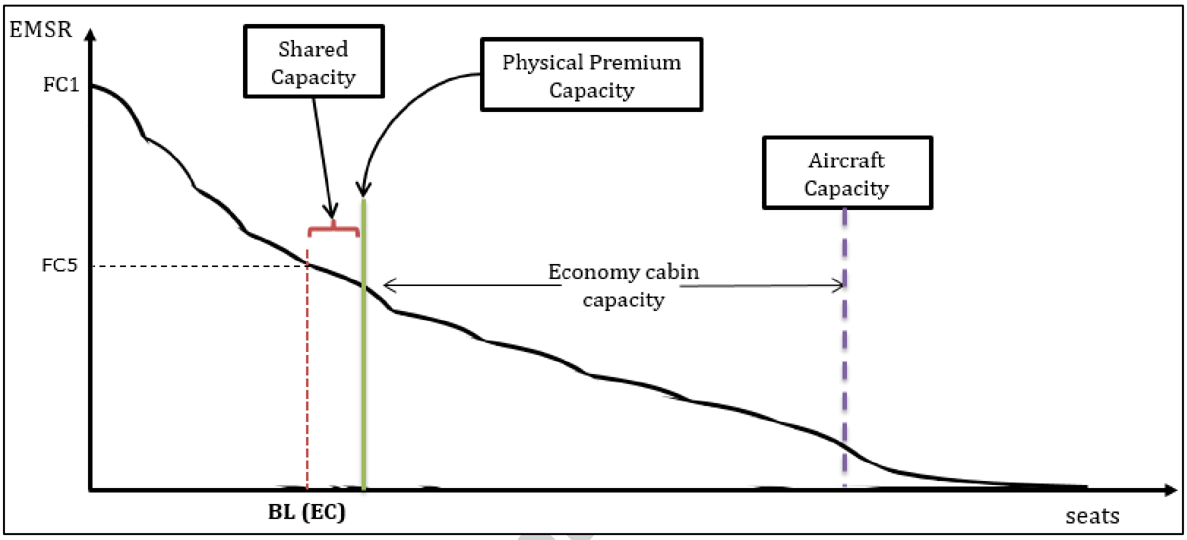

Figure 1: Full EMSR shared capacity

Table 1 provides an example of the differences between Distinct EMSRb and Full EMSR on a hypothetical leg operated by an aircraft with 20 seats in the premium cabin and 150 seats in the economy cabin. In the distinct case, the booking limit determined for FC5 is 150, corresponding to the capacity of the economy cabin; similarly, the bookings in FC1 are limited to 20 because of the premium cabin capacity. In contrast, Full EMSR applies EMSRb considering an aircraft with 170 seats, ignoring the distinction between cabins. In the example, 16 seats are jointly protected for FC1, FC2, FC3 and FC4 from the lower fare classes (FC5 to FC10), setting a booking limit of 154 for FC5. Then, the joint protections and booking limits are adjusted for the premium fare classes by subtracting the economy cabin capacity in order to reflect the real capacity of the premium cabin; e.g. the booking limit of FC4 is changed from 156 to 6. Meanwhile, the booking limits for the economy fare classes are left as calculated using Full EMSR. Although the booking limits for FC1 and FC5 are 20 and 154 respectively, it does not mean that Full EMSR allows a premium cabin seat to be sold twice (for a total number of bookings of 174). Instead, the booking limits in the premium fare classes are reduced when the bookings in economy fare classes exceed the economy cabin capacity.

		Demand Forecast		Distinct EMSRb		Full EMSR			
Fare Class	Fare	Mean	Sigma	Joint Prot.	Booking Limit	Joint Prot.	Booking Limit	**Adj. Joint Prot.**	**Adj. B.L.**
1	$1,496	3	1	2	20	2	170	**2**	**20**
2	$1,122	5	3	6	18	6	168	**6**	**18**
3	$822	8	5	14	14	14	164	**14**	**14**
4	$598	13	8		6	16	156	**20**	**6**
5	$748	16	5	14	150	43	154	**27**	**154**
6	$457	24	8	36	136	68	127	**52**	**127**
7	$363	18	5	54	114	85	102	**69**	**102**
8	$327	27	4	83	96	115	85	**99**	**85**
9	$252	35	7	119	67	152	55	**136**	**55**
10	$204	50	8		31		18		**18**

Table 1: Distinct vs Full EMSR with shared capacity

It should be noted that Full EMSR does not necessarily lead to sharing premium capacity. Joint protection for premium fare classes is limited to the premium capacity. Table 2 presents an example: the joint protection for premium fare classes calculated using Full EMSR is 24 and, correspondingly, the booking limit for FC5 is 146; however, bookings for premium fare classes cannot exceed the number of seats in the premium cabin. In this case, since the booking limit for FC5 has to be increased by 4 (from 146 to 150), the booking limits of all the economy fare classes are also adjusted by adding 4 to the value initially calculated using Full EMSR.

		Demand Forecast		Distinct EMSRb		Full EMSR			
Fare Class	Fare	Mean	Sigma	Joint Prot.	Booking Limit	Joint Prot.	Booking Limit	**Adj. Joint Prot.**	**Adj. B.L.**
1	$1,496	4	2	2	20	2	170	**2**	**20**
2	$1,122	6	3	8	18	8	168	**8**	**18**
3	$822	10	4	19	12	19	162	**19**	**12**
4	$598	11	4		1	24	151		**1**
5	$748	16	5	14	150	46	146	**42**	**150**
6	$457	24	8	36	136	70	124	**66**	**128**
7	$363	18	5	54	114	88	100	**84**	**104**
8	$327	27	4	83	96	117	82	**113**	**86**
9	$252	35	7	119	67	134	53	**130**	**57**
10	$204	50	8		31		36		**40**

Table 2: Distinct vs Full EMSR without shared capacity

3.1.2 Capacity Sharing with OD-Control RM

As in the leg-control RM case, the baseline implementation of OD-control in dual cabin aircraft considers each of the cabins distinctly; that is, a bid price for premium cabin PBP_k and a bid price for economy cabin EBP_k are calculated for each leg k to be used when determining whether a booking request is accepted. To enable the premium cabin capacity sharing capability using a Bid Price Control mechanism (such as ProBP), the "Effective Bid Price" scheme was used: a joint bid price JBP_k is calculated for each leg k considering the full remaining seating capacity; its calculation ignores the division between premium and economy cabin.

Then, the Effective Premium Bid Price for leg k, $EPBP_k$, and the Effective Economy Bid Price for leg k, $EEBP_k$, are defined as follows:

$$EPBP_k = max(PBP_k, JBP_k)$$

$$EEBP_k = min(EBP_k, JBP_k)$$

For an ODF j that traverses the set of legs L_j, this sharing mechanism makes premium cabin seats available for fare class i when the fare $F_{i,j}$ is higher than the sum of the effective bid prices corresponding to the cabin of each fare class; that is, FC's 1 to 4 are compared to Effective Premium Bid Prices and FC's 5 to 10 are compared to Effective Economy Bid Prices.

4. Passenger Origin Destination Simulator (PODS)

PODS is the simulation tool used for testing the outcome of the cabin sharing strategies proposed in this paper. This software replicates the interactions between passengers wishing to travel in different OD markets and airlines offering air transportation services in such markets. PODS allows us to evaluate the impact of the proposed strategies in a realistic environment where passengers not only choose among competitors with different flight offerings but also choose between premium and economy fare classes based on their preferences and willingness to pay. We provide a brief overview of PODS; more details of the simulation and underlying models can be found in Escovar Alvarez (2016), Fiig et al. (2010) and Tam et al. (2008).

The architecture of PODS consists of two major components: the Passenger Choice Model and the Revenue Management System. While the former generates the passenger demand, characteristics and choices, the latter replicates the process from an airline's perspective, considering historical information, forecasts and RM optimizers.

Another important characteristic of PODS is its consideration of the time dimension. For each sample (departure date), there is a booking period of 63 days during which passengers can book their itineraries. This booking period is divided into 16 time frames ("TF"), each of these representing a point in time before the departure with a specified duration.

A hypothetical network (Network V1) was designed within PODS for testing scenarios with airlines using dual cabin aircraft configurations. V1 replicates a network in which four airlines compete for passengers wishing to travel in the 572 origin-destination markets served by 442 legs that connect 44 cities (4 hubs and 40 spokes). Among those cities, 10 are considered international (or long-haul) destinations and any market that includes these destinations, either as the origin or as the destination, is categorized as an international market; meanwhile, markets connecting any two of the remaining 34 cities are considered domestic. Four airlines operate in a hub-and-spoke network structure, with a handful of point-to-point legs offered by some of the airlines. The network of Airline 1 ("AL1") is illustrated in Figure 2.

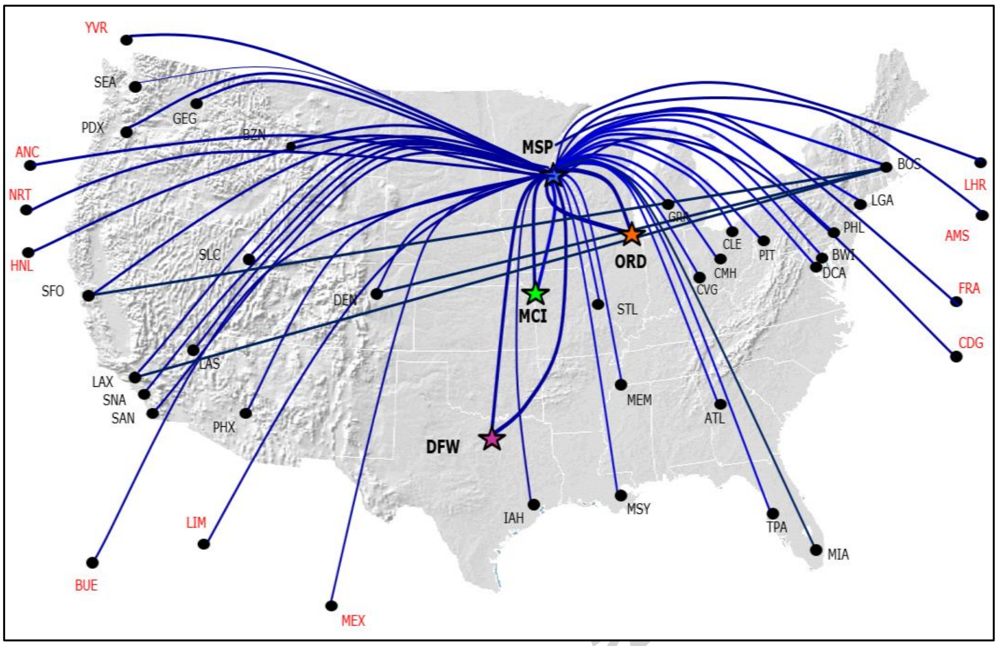

Figure 2: Airline 1 Route Map

Ten fare classes ("FCs"), each of these specifying a price and a set of conditions, are defined for each market served by the airlines in V1. These FCs are divided in two subsets: FCs 1 to 4 give access to seats in the premium cabin, FCs 5 to 10 are designated for passengers purchasing tickets in the economy cabin. Other than the obvious geographical distinction between international and domestic markets, there are also considerable differences between the fare restrictions applicable to each of them, specifically:

- AP: Advance Purchase requirement, ranging from 0 to 21 days.
- R1: Saturday night stay requirement
- R2: Cancellation or change penalty
- R3: Non-refundability

For our simulations of capacity sharing mechanisms, we also introduced additional disutilities to differentiate the premium cabin fare classes from the economy fare classes, as follows:

- D1: Disutility of Not Sitting in Premium Cabin
- D2: Disutility of Not Sitting in Premium Cabin on International Flights

Table 3 presents the sets of conditions, restrictions and disutilities applicable for international and domestic products. International products have a more restricted fare structure than domestic products. Also, the highest FC applicable for each cabin (i.e. FC1 and FC5) has fewer restrictions than any of the lower FC's in that same cabin. It should be highlighted as well that the highest economy cabin fare classes (i.e. FC5, FC6) are much less restricted than the lowest premium FC (i.e. FC4); this allows for an overlap between the prices offered for travelling in different cabins, controlled by the AP requirement that makes FC4 available only until 21 days before departure while FC5 is open until the day of departure. Table 4 provides the average fares for the international and domestic products.

		International Markets						Domestic Markets					
	Fare Class	AP	R1	R2	R3	D1	D2	AP	R1	R2	R3	D1	D2
Premium	1	0	NO	NO	NO	NO	NO	0	NO	NO	NO	NO	NO
	2	0	NO	NO	YES	NO	NO	0	NO	YES	NO	NO	NO
	3	7	YES	YES	NO	NO	NO	7	NO	NO	YES	NO	NO
	4	21	YES	YES	YES	NO	NO	21	NO	YES	YES	NO	NO
Economy	5	0	NO	NO	NO	YES	YES	0	NO	NO	NO	YES	NO
	6	3	NO	NO	YES	YES	YES	3	NO	YES	NO	YES	NO
	7	7	NO	YES	YES	YES	YES	7	NO	NO	YES	YES	NO
	8	10	YES	YES	NO	YES	YES	7	NO	YES	YES	YES	NO
	9	14	YES	NO	YES	YES	YES	14	NO	YES	YES	YES	NO
	10	21	YES	YES	YES	YES	YES	21	NO	YES	YES	YES	NO

Table 3: Network V1 -Applicable Conditions and Restrictions in International Markets

	FC1	FC2	FC3	FC4	FC5	FC6	FC7	FC8	FC9	FC10
International Markets	$2,889	$2,167	$1,589	$1,156	$1,144	$1,008	$799	$673	$574	$476
Domestic Markets	$764	$671	$502	$373	$578	$405	$310	$244	$195	$161

Table 4: Network V1 - Fare Structure Economy Fare Classes in Domestic Markets

Airline 3 does not offer FC's 1 to 4 as it does not operate aircraft with premium cabin (for example, as an LCC competitor).

5. Simulation Results

All of the experiments in this paper consist of shifting some of the conditions simulated for AL1; hence, the network and revenue management systems for Airline 2 (AL2), Airline 3 (AL3) and

Airline 4 (AL4) remain unchanged in the simulations. In addition, different baselines are set for evaluating scenarios depending if it is uses leg-based or OD-control methods. The main characteristics of the baseline RM scenarios are summarized in **Error! Reference source not found.**. PODS simulations were run for three demand scenarios – "medium demand" corresponds to approximately 80% average network load factor, "low demand" represents 73% load factor and "high demand" 86% average load factor (across both cabins and all flights in the simulated network).

	Leg-based Baseline	**O-D Control Baseline**
AL1 RM	Distinct EMSRb	Distinct ProBP
AL2 RM	Distinct ProBP	Distinct ProBP
AL3 RM	EMSRb (Single Cabin)	EMSRb (Single Cabin)
AL4 RM	Distinct ProBP	Distinct ProBP

Table 5: Baseline Summary

It should be noted that all airlines are performing distinct RM for their aircraft cabins in the baseline. Consequently, the optimization of the seat inventory of each cabin considers exclusively the FC's corresponding to its cabin (i.e. FC's 1 to 4 in premium cabin and FC's 5 to 10).

5.1 Leg-based RM Premium Cabin Capacity Sharing Schemes: Full EMSR

To evaluate the performance of Full EMSR, simulations were run with the scenario described in Table 6; the only change with respect to the baseline is that AL1 uses Full EMSR instead of distinct EMSRb. Keeping the other airlines using the same strategies as in the baseline allows us to identify the effects caused exclusively by the decision of AL1 to share the premium cabin capacity.

	Leg-based Baseline	**Leg-based Evaluated Scenario**
AL1 RM	Distinct EMSRb	Full EMSR
AL2 RM	Distinct ProBP	Distinct ProBP
AL3 RM	EMSRb (Single cabin)	EMSRb (Single cabin)
AL4 RM	Distinct ProBP	Distinct ProBP

Table 6: Full EMSR Settings – Simulation Settings

The objective of implementing Full EMSR is to share the remaining capacity in the premium cabin with additional bookings of economy fare classes; therefore, the expectation is an increase in total revenue resulting from the higher revenue captured from the economy fare classes while keeping the same revenue from premium fare classes. However, the total revenue results show either a negligible positive effect, as in the low demand scenario (increase of 0.03%), or a substantial negative impact with total revenue losses of up to -0.66% in the high demand scenario (Figure 3). When evaluated by fare class type, sharing premium cabin capacity with economy fare classes results in large percentage losses in the premium fare classes FC 1 to 4 (-5.81% to -14.59%), partially compensated by proportional gains for economy fare classes 5 to 10, ranging from 1.31% to 2.54%, depending on the demand level.

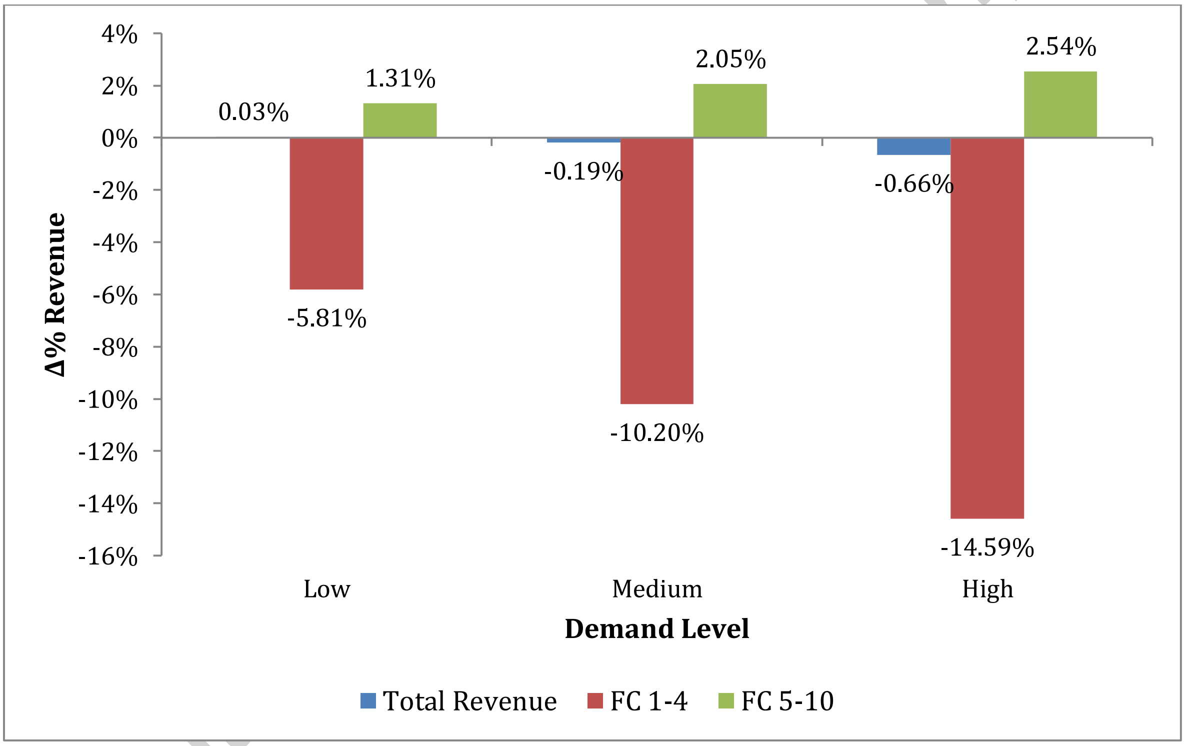

Figure 3: Full EMSR - Proportional Revenue Gains by Fare Class Type

Figure 4 shows the absolute revenue variations by fare class type and cabin. On the one hand, Full EMSR achieves the objective of increasing the revenue captured from empty capacity in the premium cabin by accommodating additional passengers that book in economy fare classes; those increases partially compensate for the significant losses in revenue from passengers that actually book in premium fare classes. On the other hand, revenue from passengers that purchase and fly in the economy cabin decreases substantially. Based on Figure 5, FC1 and FC2 are the premium fare classes that suffer the higher revenue losses, while the revenue gains come mainly from FC5 and FC10 in the economy cabin.

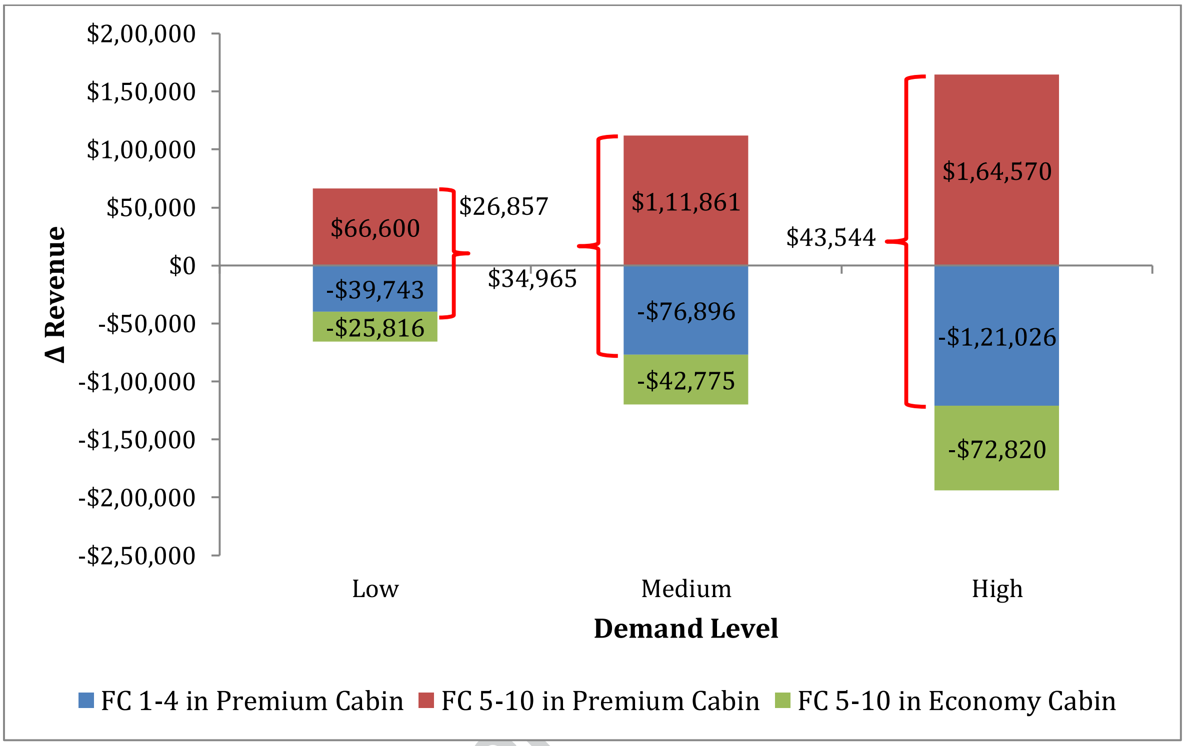

Figure 4: Full EMSR - Absolute Revenue Variation by Fare Class Type and Cabin Accommodation

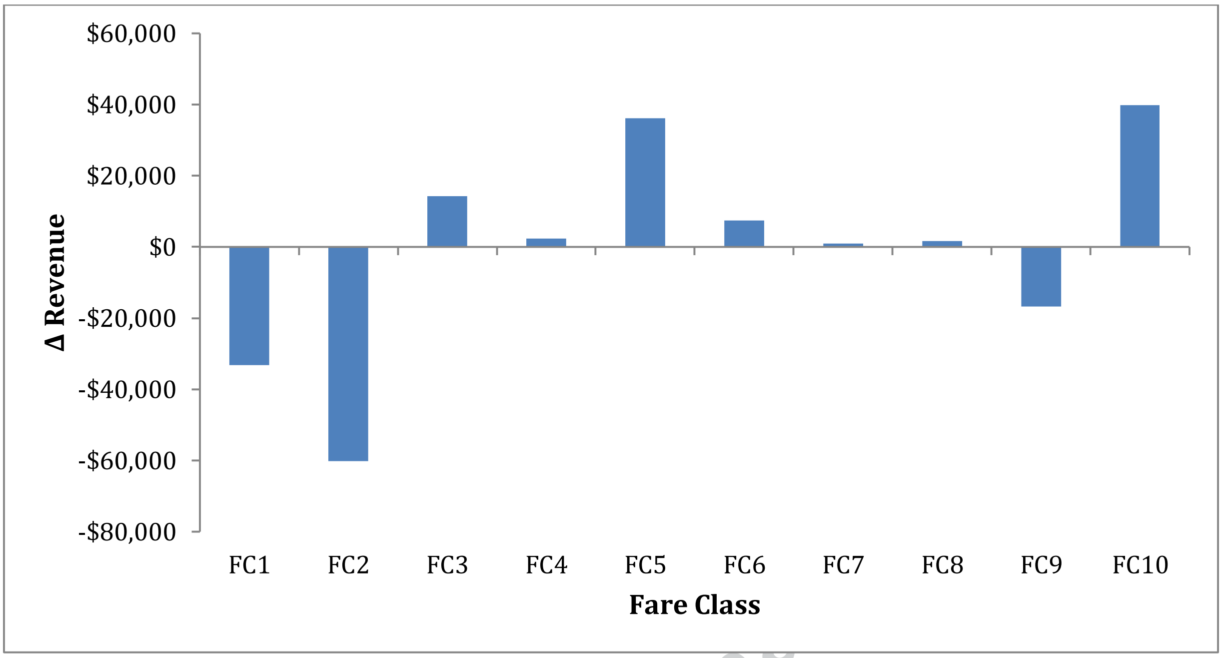

Figure 5: Full EMSR - Revenue Variation by Fare Class at Medium Demand Level

Losses in FC1 and FC2 are the result of an optimization that is fed with reduced forecasts for these fare classes. In the RM system, forecasts are based on historical observations of bookings by the airline; in this case, since some premium cabin seats are made available for economy fare class bookings, more bookings are accepted in FC5 to FC10. Given that the space in premium cabin is limited, a booking request for FC1 or FC2 might be rejected because of unavailable space. This process is similar to the "spiral down" concept, in which fewer bookings in the highest fare classes result in lower forecasts for such fare classes on future flights, reducing the number of seats to be protected, and so on, in a cyclical process. In fact, at the medium demand level, forecasts for FC1 and FC2 are 31.1% and 24.7% lower in the Full EMSR scenario than in the baseline.

It is important to note that forecast spiral down occurs despite the application of typical RM unconstraining models in the fare class forecasting process. In fact, we tested several alternative RM unconstraining models – booking curve detruncation, Expectation Maximization (EM) detruncation, as well as a substantially more aggressive Projection Detruncation approach. Although the more aggressive detruncator helped reduce forecast spiral down a little, it did not change the fundamental forecast changes in Full EMSR. Fewer high fare bookings are observed not only due to inadequate premium cabin capacity, but also with high fare class bookings lost to buy-down to more available lower fare classes and to competitors.

These changes in the forecasts have a significant effect in the protection for each of the fare classes. In the case of leg-based RM, those variations are reflected in changes to the booking limits. With the decrease in the forecasts for the highest premium fare classes, less protection from the lower premium fare classes is applied. Therefore, in the last time frames of the booking process, when most of the FC1 booking requests are expected to arrive, there is a significant reduction in availability. Rejected requests could end up booking on one of the competitors, or not flying at all.

The case of FC5 is different as additional seats are being made available for bookings in that class Therefore, considering the increase in the forecasts for FC5, more protection from the lower premium fare classes is applied, allowing the bookings in the late time frames of the booking process to increase. Finally, the nested booking limits for FC10 are also higher when premium cabin capacity is shared and explains why there is a significant increase in FC10 revenue.

5.1.1 Protection Mechanisms for Premium Fare Classes – Leg-control methods

As presented in the previous section, additional capacity made available for economy fare classes with Full EMSR results in a substantial increase in FC10 bookings, while revenue losses are observed in premium fare classes. With Full EMSR cabin sharing used throughout the booking process, the additional capacity made available to FC10 is taken by early booking low-fare economy passengers, leading to displacement of premium class passengers much closer to departure. This observation motivated the idea of a protection mechanism to delay premium cabin capacity sharing until a predetermined time frame. This aims to ensure that capacity is shared only when passengers are buying some of the high-value economy fare classes with little or no Advance Purchase (AP) requirements. Time Frame Protection ("TFP") is tested using three different predetermined time frames: TFP=10 (cabin sharing only within 14 days before departure), TFP=12 (7 days before departure) and TFP=14 (3 days before departure).

For all these cases, the highest total revenue gains occur when premium cabin capacity is made available for passengers with bookings in economy fare classes just three days before departure (TFP=14). While in the low demand scenario the gains increase from 0.03% to 0.51%, in the medium and high demand scenarios TFP turns losses into gains of 0.77% and 1.06%, respectively; Figure 6 presents the results for the medium demand scenario.

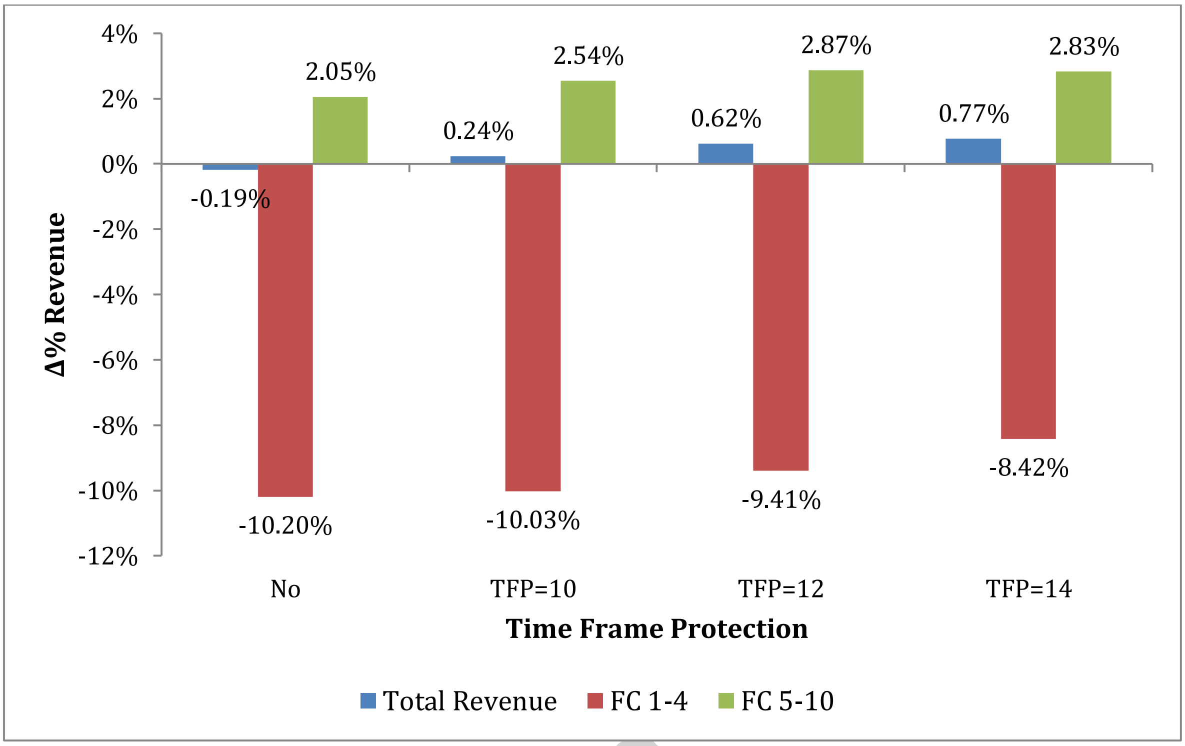

Figure 6: Time Frame Protection - Full EMSR - Medium Demand

To understand the effect of TFP on AL1's revenues when using Full EMSR it is useful to analyze how the bookings are allocated between the cabins. Figure 7 illustrates how, as expected, both the revenue gains of economy fare classes flying in premium cabin and the losses of premium fare classes decrease as premium capacity sharing is delayed. However, even the use of TFP does not eliminate the revenue losses of premium fare bookings, suggesting that the variation in the forecasts is still contributing to a spiral down effect. In fact, the total revenue captured from passengers seated in the premium cabin (regardless of the fare class type they have bought) may even decrease, as in the case of TFP=14.

Still, substantial revenue gains are captured from passengers buying economy fare classes and seated in the economy cabin; in Figure 7, for example, this category shifts from losing $42,275 to gaining $34,192 in the medium demand scenario with TFP=14. In addition, Figure 8 shows how delaying premium cabin capacity sharing reduces the losses in FC1 and FC2 and increases the gains in FC5 while reducing FC10 revenue; this indicates that available premium cabin capacity seats are shared with passengers booking in the higher economy fare classes instead of opening space for low fare class bookings, mostly FC10. However, even use of TFP=14 (3 days to departure) does not completely eliminate the revenue losses of premium fare passengers, indicating that forecast spiral down is still an issue.

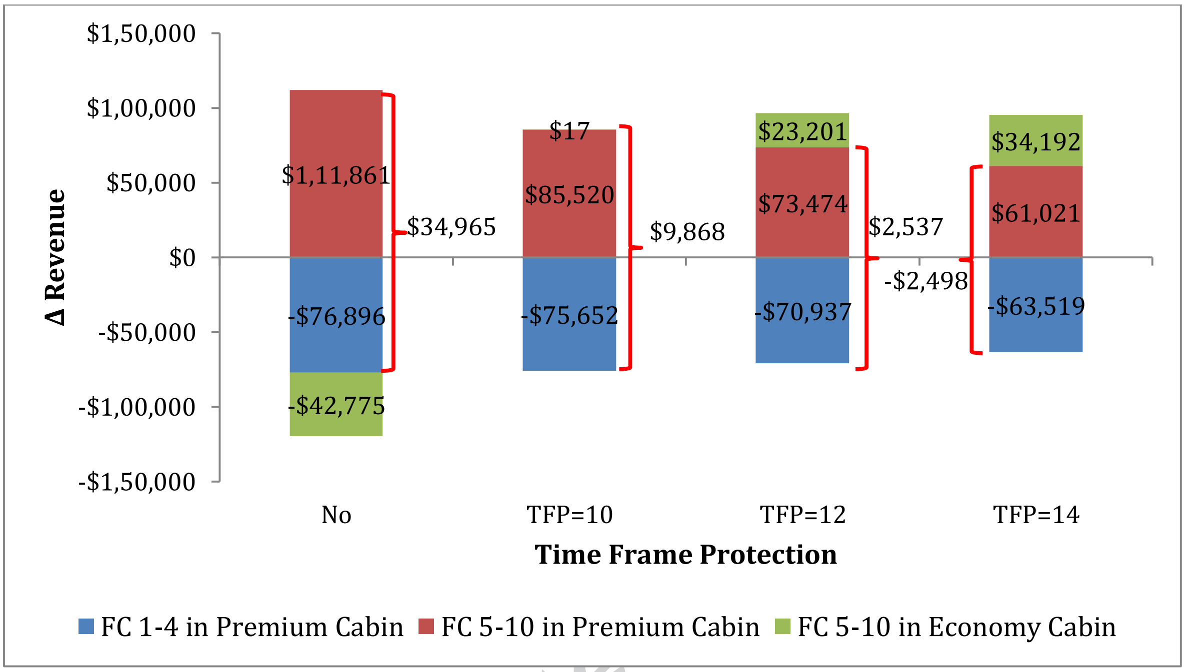

Figure 7: Full EMSR - Absolute Revenue Variation by FC Type, Cabin Accommodation and TFP – Medium Demand

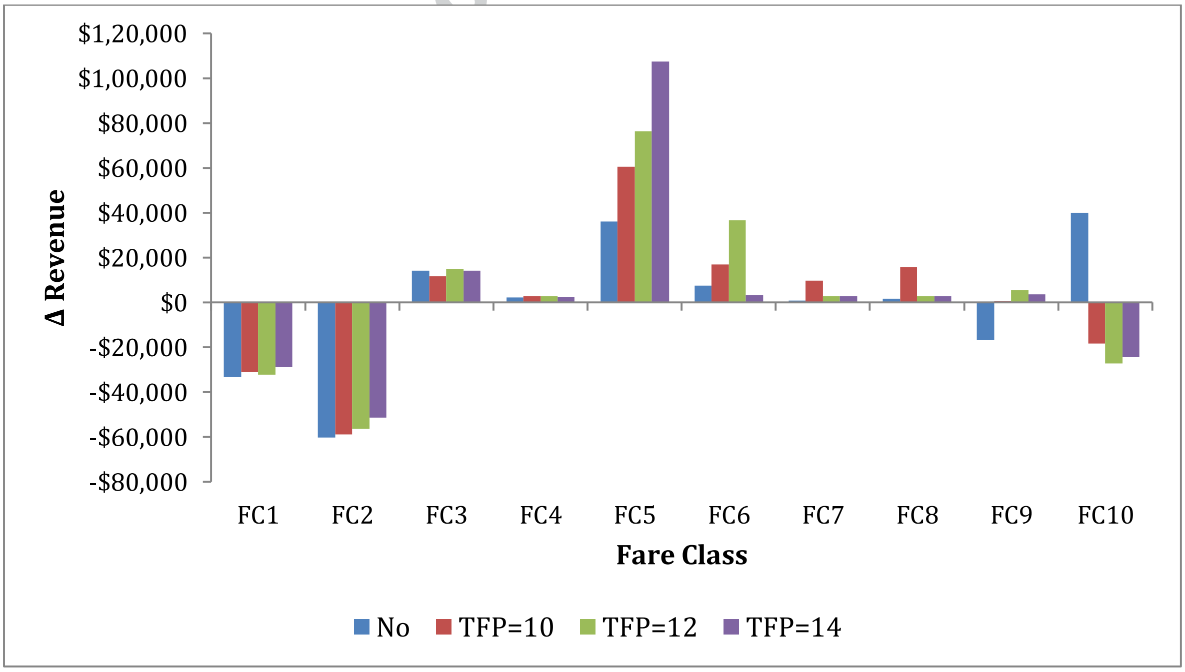

Figure 8: Absolute revenue variations by fare class – Full EMSR with TFP – Medium Demand

5.2 OD-Control RM Premium Cabin Shared Capacity: Effective Bid Price Capacity Sharing

The main settings for the simulations used to evaluate the Effective Bid Price (EBP) capacity sharing method are presented in Table 7. As was the case with full EMSR, Effective Bid Price capacity sharing is applied from the beginning of the booking period

	OD-Control Baseline	OD-Control Evaluated Scenario
AL1 RM	Distinct ProBP	Shared ProBP Effective Bid Price Sharing (EBP)
AL2 RM	Distinct ProBP	Distinct ProBP
AL3 RM	EMSRb (Single cabin)	EMSRb (Single cabin)
AL4 RM	Distinct ProBP	Distinct ProBP

Table 7: Effective Bid Price Capacity Sharing - Simulation settings

As illustrated in Figure 9**Error! Reference source not found.**, a negligible increase in AL1's total revenue (0.06%) is obtained at the low demand level; the situation worsens at the medium and high demand scenarios, where AL1 reduces its total revenue by 0.33% and 0.83%, respectively. Like in the leg-based RM scenarios, the total revenue variation consists of the combination of a considerable proportional decrease in premium fare classes revenues (ranging from -5.23% to -12.44%) with a lower proportional increase in economy fare classes revenues (1.26% to 1.86%). However, because of the larger number of seats in economy cabin, the magnitude of the absolute losses in FC's 1 to 4 is larger than the gains generated by providing additional capacity to FC's 5 to 10.

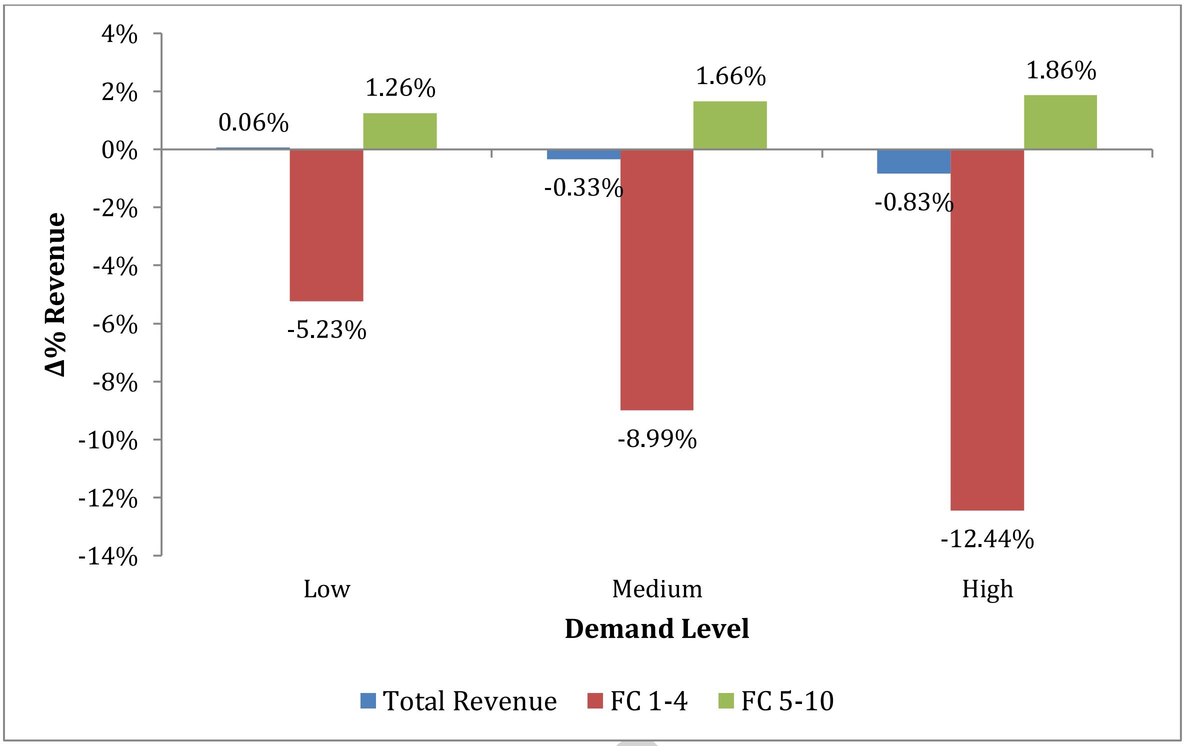

Figure 9: Effective Bid Price - Proportional Revenue Gains by Fare Class Type

EBP captures a significant amount of revenue from passengers that booked in economy fare classes by allowing them to sit in the premium cabin, compensating for the losses in the revenue from premium fare classes; overall, the revenue captured from passengers flying in premium cabin increases, as would be initially expected. However, the revenue from passengers flying in economy cabin has a dramatic decrease that exceeds the premium cabin revenue gains, as presented in Figure 10.

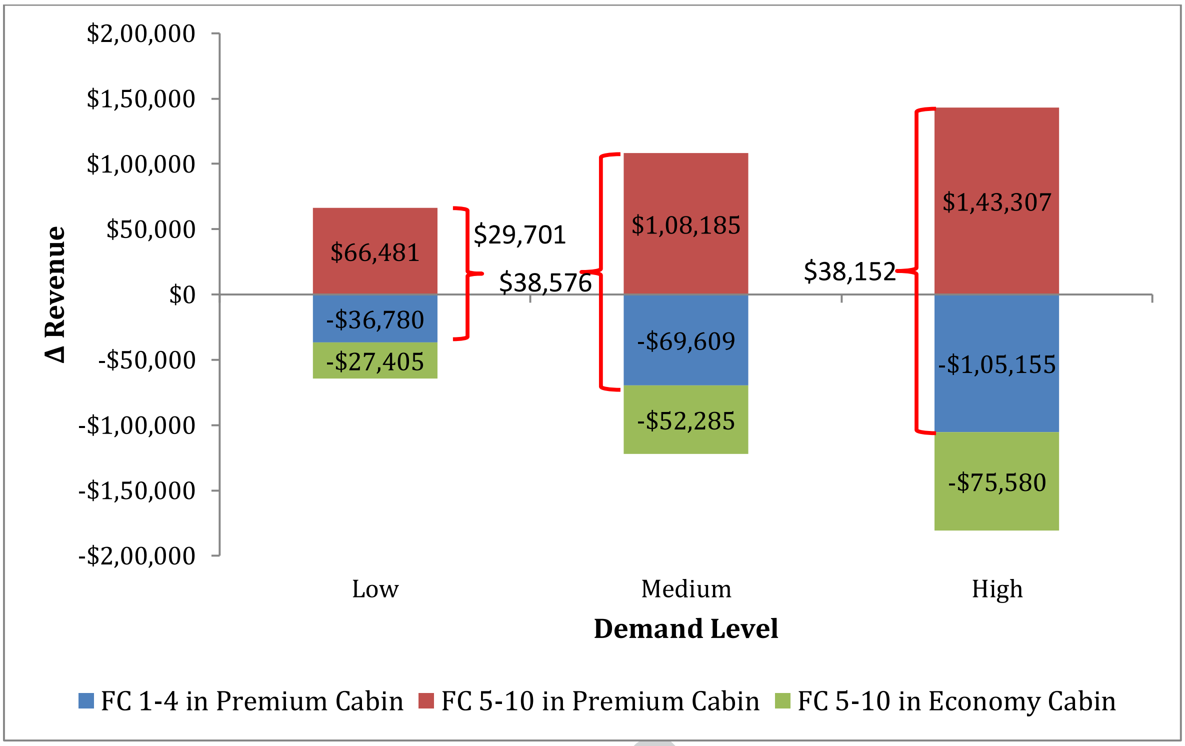

Figure 10: Effective Bid Price - Absolute Revenue Variation by Fare Class Type and Cabin Accommodation

Similar to what was found in the Full EMSR simulations (without Time Frame Protection), total revenue variations are driven mainly by substantial losses in FC1 and FC2 that are larger in absolute value than the gains in FC5 and FC10. Again, this is the result of the variation in the historical bookings recorded in each fare class over repeated samples (departure dates) and the effect of such historical bookings on the bid prices applicable for each fare class group. As illustrated in Figure 11, the premium cabin and joint cabin bid prices applicable for the EBP scenario are, on average, lower than the average premium cabin bid prices when cabins are managed distinctly during most of the booking period, with the exception of the last three days of the booking process. On the other hand, Figure 12 presents the average economy cabin and joint cabin bid prices applicable for EBP compared to the distinct economy cabin bid price of the distinct scenario. In this case, the average economy cabin bid price in the EBP scenario is lower during the early stages of the booking period (until 26 days before departure) and in the last days before departure.

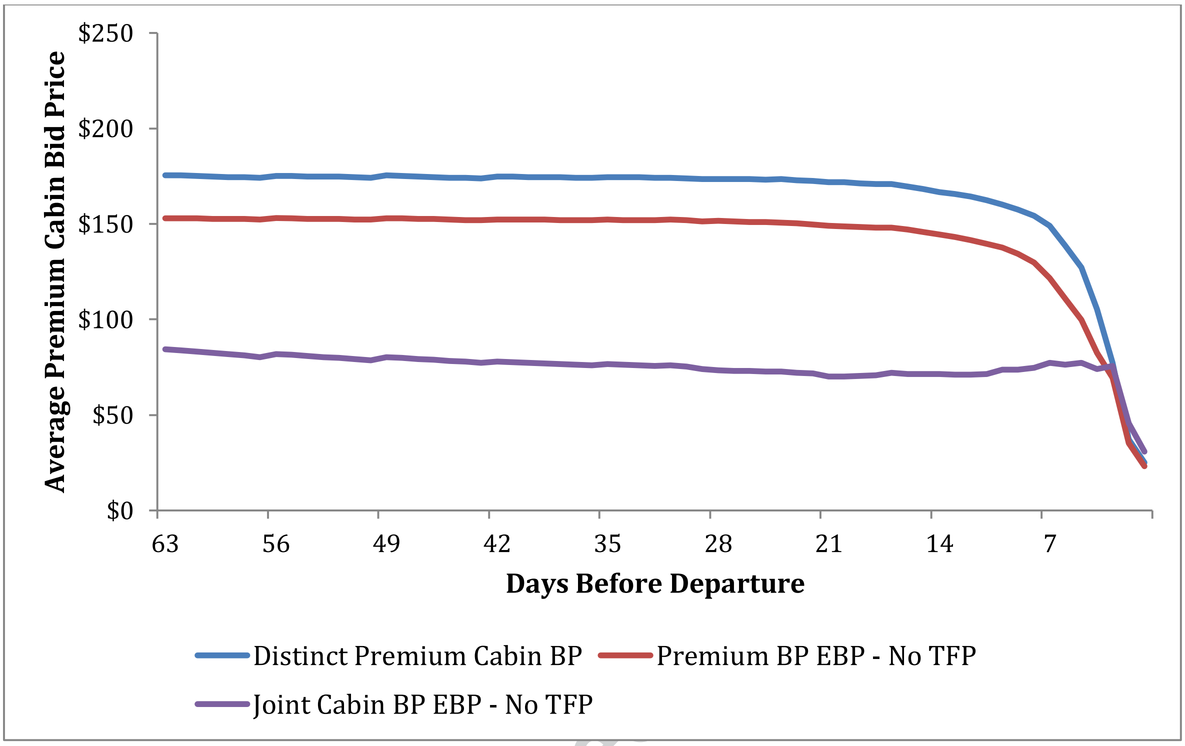

Figure 11: Premium and Joint Cabin Bid Prices – Medium Demand

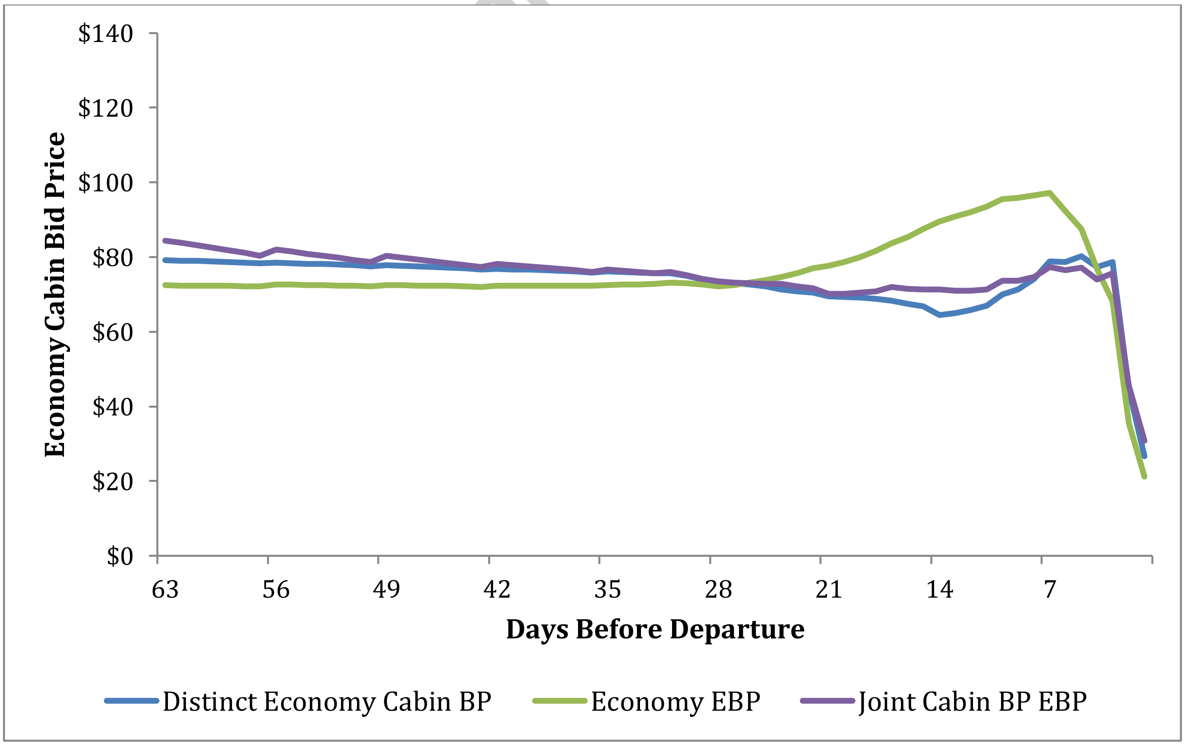

Figure 12: Economy and Joint Cabin Bid Prices – Medium Demand

During the initial time frames of the booking period the average economy cabin bid price is lower when premium cabin capacity is shared; hence, the likelihood of accepting a low fare booking request increases. Such increase reduces the available seats in economy cabin for higher economy fare classes and, as a consequence, increases the demand for premium cabin seats. In addition, a lower premium cabin effective bid price in the early stages of the booking period allows more passengers to book in the lower premium fare classes (FC3 and FC4). These interactions result in a substantial decrease in availability for the higher fare classes in the last days before departures, as shown in the closure rates for FC1.

Figure 13 presents the variation in cumulative total revenue at each demand level. In the initial time frames, the revenue captured using EBP is higher than in the baseline mostly because of the increase in FC10 bookings. However, it also shows how in the last days before departure and despite the additional bookings recorded in FC5, all the gains are lost because of the reduced availability for FC1 and FC2 bookings.

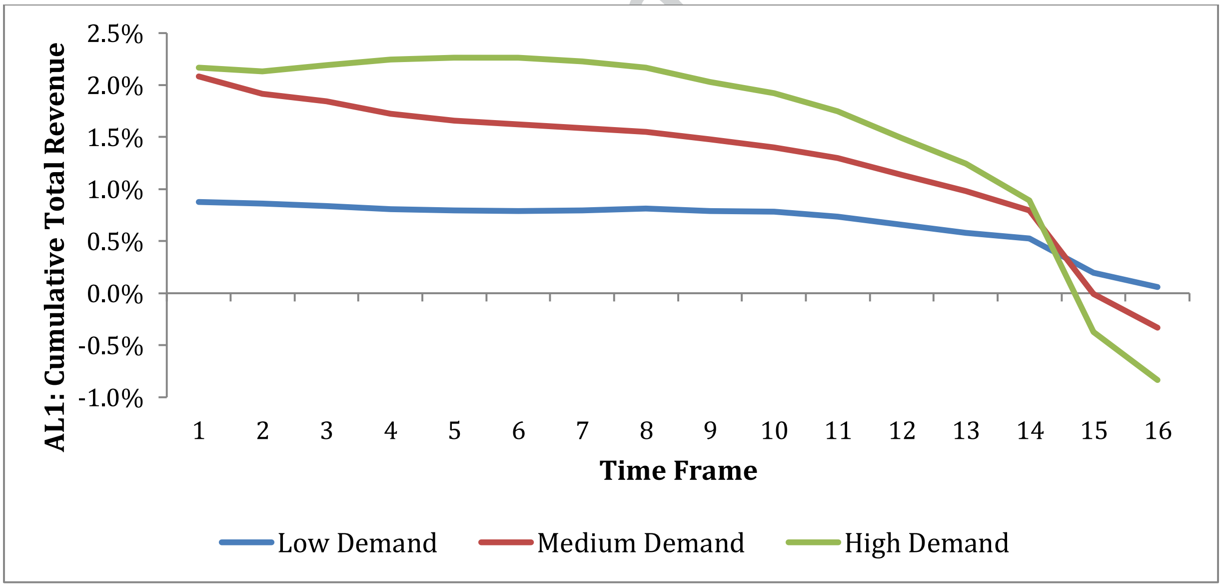

Figure 13: Effective Bid Price - Cumulative Total Revenue Variation

5.2.1 Protection Mechanisms for Premium Fare Classes – OD control methods

Just as in the case of Full EMSR in the leg-control scenarios, we found that TFP is a key element to turn losses into gains when implementing capacity sharing mechanisms such as Effective Bid Price under OD Control RM. Substantial total revenue gains are obtained when capacity is shared only at the later stages of the booking process (Figure 14). In fact, at the high demand level scenarios total revenue increases by 1.1% by delaying capacity sharing until TFP=14 (3 days to departure). Again, since the economy fare classes' revenue is substantially larger than the revenue from premium fare classes for AL1, a relatively low proportional increase in economy fare classes' revenue can compensate for much larger proportional decreases in revenue. Figure 15 shows that the main source of gains in total revenue comes from passengers flying in the economy cabin. However, TFP does not eliminate the losses in revenue from premium fare class passengers caused by the reduction in the forecasts for such fare classes.

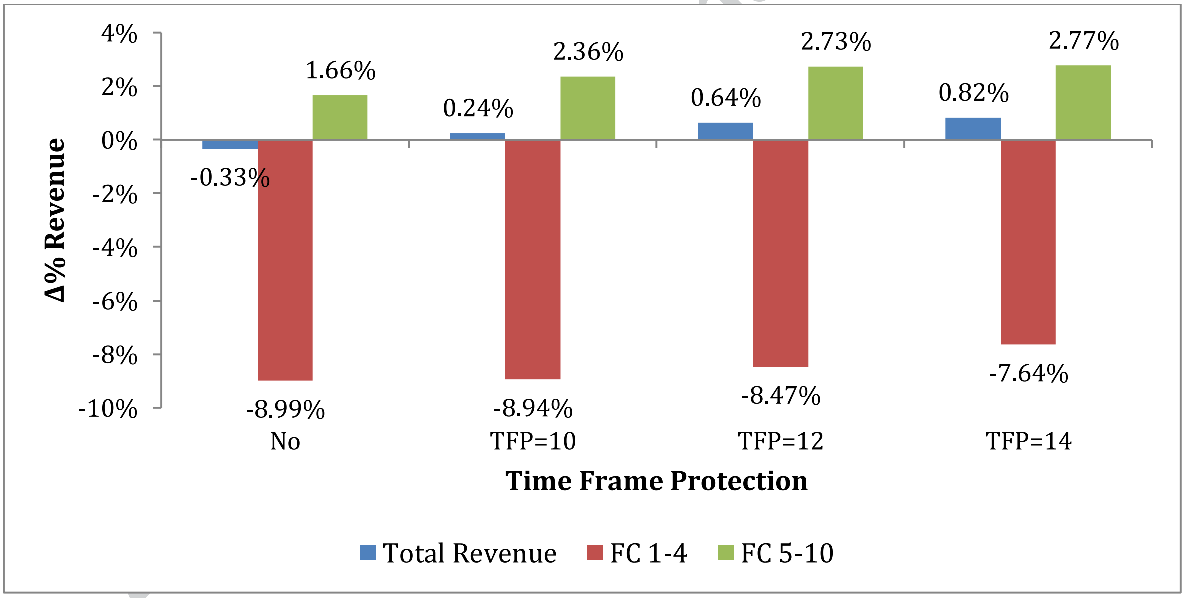

Figure 14: Effective Bid Price - Time Frame Protection - Medium Demand

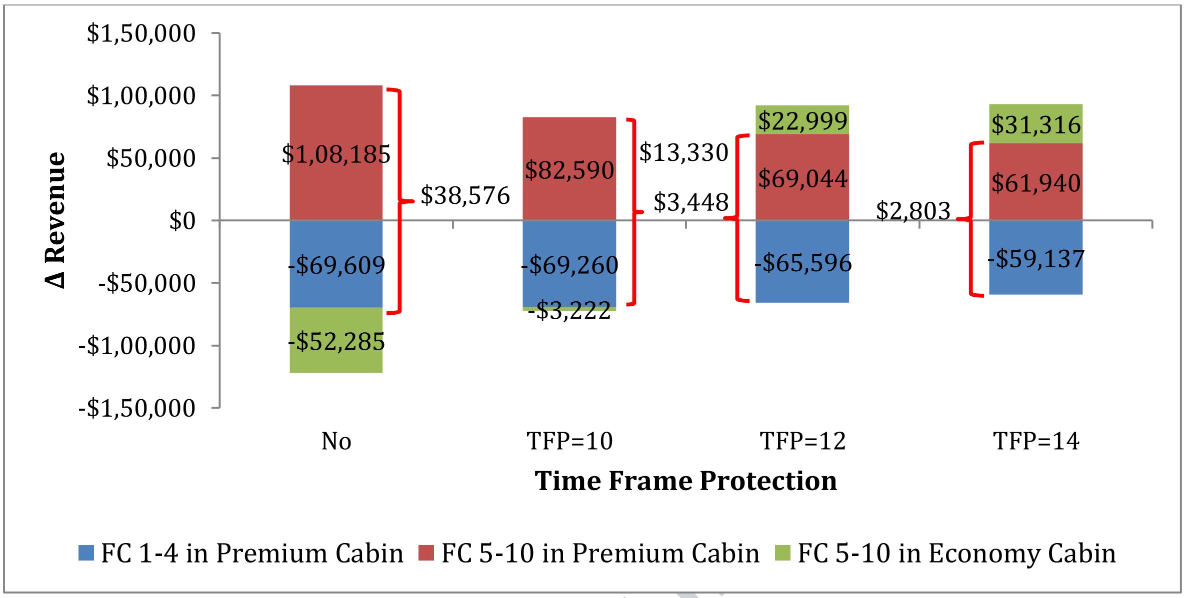

Figure 15: Effective Bid Price - Absolute Revenue Variation by FC Type, Cabin Accommodation and TFP – Medium Demand

6. Conclusions

To reduce the impact on revenues of the constraints imposed by the division of the aircraft into fixed cabins, premium cabin capacity sharing schemes were proposed for taking advantage of available premium cabin seats to accommodate additional passengers with bookings in economy fare classes. The expectation is to increase the revenue generated from the premium cabin by adding passengers with bookings in economy fare classes to the passengers that would have been expected to buy premium fare classes (fare classes 1 to 4) when cabins are managed distinctly. Full EMSR is proposed as the premium cabin capacity sharing mechanism for leg-based RM, and Effective Bid Price capacity sharing is designed for OD-control RM.

The results of the simulations in PODS show that, if capacity is shared from the beginning of the booking period, both mechanisms result either in a decrease or a negligible increase in total flight revenue (depending on the demand level). In these cases, although the revenue captured from economy fare classes increases because of the additional capacity made available, the absolute revenue losses associated with premium fare classes are larger. Premium cabin capacity sharing without implementing an additional protection mechanism allows many low-value economy fare class bookings to be accepted in the initial time frames, resulting in the rejection of a larger proportion of high value premium fare class passengers closer to departure. As the historical database records more bookings in economy fare classes and less bookings in premium fare classes, the forecasts used for determining the protection levels for the highest fare classes are affected and, as a consequence, availability for low fare classes increases (similar to the spiral down concept).

Time Frame Protection is proposed as a mechanism to delay the activation of premium cabin capacity sharing until after a predetermined time frame in the booking process. As presented in Table 8, total revenues increase when sharing is activated just in the last frames of the booking process; such revenue increases are achieved by increasing the revenue gains from economy fare classes and decreasing the losses from premium fare classes. More specifically, the best results are observed when premium cabin capacity is shared very late in the booking process, such as just three days before departure (TFP=14).

		Full EMSR			**Effective Bid Price**		
	Demand Level	Low	Medium	High	Low	Medium	High
No protection	Total	0.03%	-0.19%	-0.66%	0.06%	-0.33%	-0.83%
	FC1-FC4	-5.81%	-10.20%	-14.59%	-5.23%	-8.99%	-12.44%
	FC5-FC10	1.31%	2.05%	2.54%	1.26%	1.66%	1.86%
TFP=14 (3 days before departure)	Total	0.51%	0.77%	1.06%	0.58%	0.82%	1.10%
	FC1-FC4	-4.96%	-8.42%	-11.70%	-4.53%	-7.64%	-10.34%
	FC5-FC10	1.72%	2.83%	3.98%	1.73%	2.77%	3.76%

Table 8: Summary of proportional revenue variations for Full EMSR and Effective Bid Price with and without Time Frame Protection

Finally, Table 9 compares what an airline might intuitively expect with respect to premium cabin capacity sharing against the results found in the simulations. First, it is found that although the intention of an airline is to increase its total revenue, this is only achieved when capacity is shared late in the booking process. Another expectation is to increase revenue from passengers sitting in the premium cabin by keeping the original revenue coming from passengers booking in premium fare classes and adding the revenue from passengers booking economy fare classes that can use empty premium cabin seats. However, in all cases we observed a reduction in the revenue from passengers booking in premium fare classes (FC 1-4). Even in the scenarios in which capacity is shared only in the last days before departure the total premium cabin revenue (passengers ultimately sitting in the premium cabin) either decreases or just increases slightly. Similarly, the expectation of capacity sharing might be that total economy cabin revenue (passengers flying in economy cabin) remains relatively unchanged. However, with capacity sharing and no Time Frame Protection, total economy cabin revenue decreases due to the shift in booking mix to lower fare classes. This shift is ameliorated with TFP=14.

	Expected	Simulations without Time Frame Protection	Simulations with TFP=14
Total Revenue (premium + economy)	Increases	Increases slightly/Decreases	Increases
FC 1-4 Revenue	Unchanged	Decreases	Decreases
FC 5-10 Revenue	Increases	Increases	Increases
Premium Cabin Revenue	Increases	Increases	Increases slightly/Decreases Slightly
Economy Cabin Revenue	Unchanged	Decreases	Increases

Table 9: Comparison of Expected Results and Simulations Results

6.1 Suggestions for Future Research

In this paper Probabilistic Bid Price Control ("ProBP") is the OD-control method used for maximizing the revenue of an airline at the network level. However, there are other OD-control mechanisms such as Displacement Adjusted Virtual Nesting (DAVN) and Unbucketed Dynamic Programing (UDP). Exploring the performance of those methods applied to dual cabin RM and understanding the differences in the results is an alternative that can potentially to better results. Similarly, other premium cabin capacity sharing mechanisms and protection schemes could be proposed, focusing mainly on keeping the total revenue gains at least at the levels achieved with the methods proposed herein while decreasing the losses in the premium fare classes.

The results of this study indicate that sharing premium cabin capacity leads to a process that is similar to the "spiral down" concept; that is, reduced forecasts for premium fare classes and increased forecasts for economy fare classes result in lower protection levels for the former and additional capacity made available for the latter. One of the methodologies typically used in revenue management for reducing the impact of spiral down is hybrid forecasting; this mechanism considers together product-sensitive and price-sensitive demands by incorporating concepts of willingness to pay (Belobaba & Hopperstad, 2004). Hybrid forecasting is typically paired with a fare adjustment scheme that uses sell-up estimates to estimate the total demand available in each fare class if such fare class was the lowest open (Fiig et al, 2010). Therefore, the evaluation of the impact of hybrid forecasting and fare adjustment in dual cabin aircraft revenue management represents an interesting opportunity for research.

Finally, as described in Section 3, this study does not consider upgrades associated with frequent flyer programs or sales of paid upgrades. Assessing the impact of premium cabin capacity sharing on those upgrade mechanisms (and vice versa) could open a new avenue of future research.

References

[1] Alstrup, J., Boas, S., B.G. Madsen O., Vidal R. (1986), *Booking policy for flights with two types of passengers.* European Journal of Operational Research 27, pp. 274-288.

[2] Belobaba, P., and Hopperstad, C. (2004), *Algorithms for Revenue Management in Unrestricted Fare Markets*. Cambridge, MA: Presentation to Informs Revenue Management Section Meeting.

[3] Escovar Álvarez, G. (2016) *Capacity Management Schemes for Dual Cabin Aircraft : Airline Revenue Management Insights*. Cambridge, MA: Master's Thesis, Massachusetts Institute of Technology.

[4] 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.

[5] Gallego, G., and Stefanescu, C. (2009) *Upgrades, Upsells and Pricing in Revenue Management*. Working Paper. Available at SSRN: http://ssrn.com/abstract=1334341

[6] Lepage, P. O. (2013). *Performance of Multiple Cabin Optimization Methods in Airline Revenue Management*. Cambridge, MA: Master's Thesis, Massachusetts Institute of Technology.

[7] Rauch, J., and Poelt, S. (2015). *Upgrading and Capacity Sharing – Dual Cabin Method for PODS*, Lufthansa German Airlines.

[8] Tam, W., Belobaba, P., Hopperstad, C. (2008) *Passenger Origin-Destination Simulator (PODS) Summary of Processes and Models*. Unpublished Report, Massachusetts Institute of Technology

[9] Walczak, D., (2010) Capacity sharing. In *50th AGIFORS Annual Proceedings*. Curran Associates, Red Hook, NY.