Capacity Management Schemes for Dual Cabin Aircraft: Airline Revenue Management Insights

by

Germán Escovar Álvarez
B.S. Industrial Engineering
Universidad de los Andes, Colombia (2007)

Submitted to the Department of Civil and Environmental Engineering in partial Fulfillment of the Requirements for the Degree of

MASTER OF SCIENCE IN TRANSPORTATION

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2016

Author……
Department of Civil and Environmental Engineering
May 13, 2016

Certified by……
Peter P. Belobaba
Principal Research Scientist of Aeronautics and Astronautics
Thesis Supervisor

Accepted by……
Heidi Nepf
Donald and Martha Harleman Professor of Civil and Environmental Engineering
Chair, Graduate Program Committee

Capacity Management Schemes for Dual Cabin Aircraft: Airline Revenue Management Insights

by

Germán Escovar Álvarez

Submitted to the Department of Civil and Environmental Engineering
on May 13, 2016, in Partial Fulfillment of the Requirements for the Degree of
Master of Science in Transportation

Abstract

The selection of an aircraft type has long term effects on the competitive position of the airline. In a market that is continuously evolving, such as the commercial aviation industry, any degree of flexibility for adjusting an aircraft capacity to match better the demand provides an opportunity for airlines to remain profitable when conditions have changed.

This thesis focuses on airlines operating dual cabin aircraft (premium and economy cabins) and explores two alternatives that can be used to adjust the capacity made available to maximize revenues. On the one hand, an easily implementable strategy of premium cabin capacity sharing is proposed with the intention of allowing passengers booking in economy fare classes to be accommodated in premium cabin seats when these seats are expected to be empty. On the other hand, a medium to long-term solution of changing the aircraft configuration (through aircraft replacement or retrofit) is considered. Both alternatives are tested using simulation tools that incorporate revenue management concepts and passenger decision making.

Four heuristics are developed and tested to evaluate premium cabin capacity sharing. Based on the simulations, it is found that the methodologies proposed can generate total revenue gains of up to 1.1%. Nevertheless, 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 could result in revenue losses for the airline.

With respect to cabin configuration analysis, an analytical model based on the Boeing-Swan Spill Model (BSM) is applied to dual cabin aircraft and is used to estimate the impacts on revenue due to a change in configuration. These results are compared to the results of the simulations and it is found that the BSM is able to predict in most cases whether the configuration change will generate revenue gains or losses for the airline. However, estimates of the dual cabin BSM ignore the interaction between passengers of both cabins, leading to incorrect estimates of load factors and average revenue values of spilled or accommodated passengers.

Thesis Supervisor: Peter P. Belobaba
Title: Principal Research Scientist of Aeronautics and Astronautics

Acknowledgments

I would like to express my gratitude to Professor Peter Belobaba for his guidance, advice and support over the last two years. It was a privilege to learn, work, and exchange ideas with an expert in the field. His emphasis on time management (and deadlines!) and his dedication of time were key elements for the success of this thesis. Moreover, his decision of offering me a Research Assistantship that allowed funding my studies was essential to make my MIT experience possible.

I would also like to thank Craig Hopperstad not only for his help developing the PODS capabilities required for my thesis but also for his responsiveness, patience and good attitude.

I thank Mr. Bockelie (Adam), Mr. Bachwich (Alex), Salz (Alex), Daniel George Ambrose, *Monsieur* de Vergnes (Matthieu), Mr. Wittman (Mike) and Mr. Petraru (Oren) for being a great team to work with. They made it easier for me to survive the PODS conferences and helped to make our trips to Montreal, Atlanta, Houston and Dallas enjoyable. It was also great that we were able to keep our lab working 24/7 (as an airline!) by combining our "shifts". In addition, I would like to thank my MST friends Javier, Alice, Nate, Katie, Ryota (Tommy) and Patton for the great moments I have shared with them.

Finally, I thank my parents, my girlfriend and my family for their unconditional love and support throughout these years.

Table of Contents

List of Figures

List of Tables

1. Introduction

The operation of a flight by a commercial airline is the result of many decisions taken in different stages of its planning process. In the short to medium term, tactical decisions related to the operation of a flight, such as crew scheduling, maintenance actions and pricing and seat inventory management, are required. In the medium to long term, strategic planning, marketing strategy, commercial alliances, information technology and fleet and route planning are some of the main topics on the agenda of an airline management team. Among these many complexities, there are two dimensions that are particularly relevant to this thesis: on the one hand, the role of differential pricing and revenue management as tools that have been developed to maximize the revenue obtained from the available seat capacity. On the other hand, the definition of the number of seats installed in each of the cabins (premium and economy) of an aircraft affects the total number of seats that can be offered to the travelers. When combined, both dimensions can have a substantial impact on the profitability and competitive position of the airline.

1.1 Differential Pricing and Revenue Management

Airlines are often studied in transportation economics courses as one of the most successful cases of the implementation of differential pricing as a mechanism to increase the producer surplus (through the increased capture of consumer surplus). Differential pricing consists mainly of two different tools: price discrimination and product differentiation. The former consists of charging different prices for the same product based in the estimated willingness to pay of different groups of customers; by contrast, product differentiation refers to the practice of assigning prices to the products based on their characteristics (i.e. higher prices for products with higher value to customers or higher production costs for the airline).

The main challenge associated with this process is that rational travelers would always prefer to pay less for their trips. Therefore, airlines need to find mechanisms that effectively stimulate passengers to buy tickets at a price that is closer to their willingness to pay, instead of selecting a lower fare. The first instrument that airlines have developed to achieve this goal are fare restrictions; some of these restrictions are associated with charging fees for potential changes or cancellations while others are related to requirements such as advanced purchase, minimum stay or round-trip travel. The combination of these restrictions (and different levels of each) leads to a wide variety of fare products. Depending on the specific characteristics of each passenger, the total inconvenience associated with each of these combinations would be perceived as an increased cost of the fare product being purchased. Based on such increased disutilities, a passenger might (or not) prefer to buy a more expensive, but less restricted, fare product to fly between two cities.

The second instrument available for airlines to offer different prices is the provision of amenities, in addition to the service of transportation from A to B, that have a direct effect on the utility perceived by the passenger. Some examples of these added value items include the option of carrying checked bags at no additional cost, additional frequent flyer program points, lounge access, pre-assigned seats and preferential check-in lines. Moreover, airlines have the option of dividing their aircraft into cabins; in one or more of these cabins, an enhanced service is offered to satisfy the expectations of passengers with high willingness to pay. The services provided in such cabins often include premium food and drinks, additional space, more comfortable seats and service oriented flight attendants. While Business Class service is the most typical premium cabin service and is provided by many carriers (mostly international), First Class is used exclusively by some world-class carriers; moreover, a small group of airlines has developed cabins that exceed the service usually provided in First Classes through the installation of suites featuring elements such as beds and showers. Another example is the recent introduction of "premium economy" seats, featuring seats that are usually comparable to those in economy cabin, but providing more legroom and other perks (such as free drinks and snacks) to the passengers.

Once products (and their associated prices) are defined by an airline, airlines use revenue management techniques to maximize the revenue captured from the flights in its network. This means managing the seat inventory of the aircraft such that each seat captures as much revenue as possible given the characteristics of the demand for each market. This is achieved by protecting seats for travelers that arrive late in the booking process but are willing to pay some of the highest fares available. As discussed herein, dividing the aircraft in cabins with different service standards helps the airline to incentivize passengers with high willingness to buy high fare products. However, dividing the aircraft in cabins adds an additional restriction to the process of revenue management, since it limits the amount of seats to be protected for the higher fare classes from the lower fare classes to the physical capacity of the premium cabins. In addition, when cabins are managed distinctly, airlines cannot use empty premium cabin seats for accommodating passengers that are willing to buy some of the lower fare classes.

1.2 Fleet Planning and Aircraft Configuration Process

Fleet planning is one of the long term decisions faced by an airline in its planning process. Its main objective is to determine the number and type of aircraft required for serving the proposed business plan of the airline for any period in the future. Notwithstanding the approach used for making these choices, incorporating an aircraft type in an airline's fleet is typically made with a time horizon of at least 10 years of operation. In addition, because of the availability of future slots and the specificities associated with aircraft manufacturing, purchase agreements between aircraft manufacturers and airlines usually involve aircraft to be received four or more years in the future. This means that airlines make fleet decisions that will impact their performance many years ahead of such decisions; as expected, there is high uncertainty about the market conditions prevailing at the time when selected aircraft will be operated.

Once an aircraft type has been chosen, airlines have some other decisions to be made related with the specificities of such airplane. Together with the selection of the engine type to be installed (when alternatives are available), the other major process consists in choosing (or designing, if applicable for the aircraft type) a layout of passenger accommodation ("LOPA") where the details of the aircraft's cabin interior are defined. If the option is available, airlines can decide either to stick to the number of seats used as part of the assumptions of the fleet planning process or to proceed with other configuration. Since the final decision is made a couple of years before the delivery of the first aircraft of an order, this option provides the option to the airline of making an adjustment if the characteristics of the markets to be operated by such aircraft type have changed. Furthermore, cabin configuration decisions are not limited to the fleet planning process, as retrofit programs can be undertaken with existing aircraft in the fleet (Nita and Scholz, 2009).

Despite this degree of flexibility, the decision of how many seats to install in each of the cabins is far from simple, as it has implications for the revenues and costs of the airline, as well as in other items that are more difficult to quantify. It should be considered that in order to provide more space and comfort to each of the passengers flying in premium cabin, such seats are typically wider than economy cabin seats and increased pitch is allowed between rows. Therefore, depending on the characteristics of the aircraft type and of the seats, adding a premium cabin seat represents decreasing more than one seat in economy cabin. An example of this trade-off is presented in Figure 1-1Figure 1-1: Major US airline Boeing 767 and Boeing 777 seats by aircraft adapted from the website of a major US airline, for their Boeing 767 and Boeing 777-200 aircraft. In the case of the Boeing 767, the elimination of its First Class (10 seats) plus the removal of 2 Business Class seats from the Low Density configuration allows for the installation of 25 Main Cabin Extra (premium economy) seats and 28 additional Main Cabin seats, resulting in a total of 41 additional seats in the High Density configuration. For the case of the Boeing 777-200, the High Density configuration is able to carry a total of 13 additional passengers as a result of the change of its configuration. Another case is found in the Airbus A320 aircraft characteristics documents, where two different A320 configurations are presented: the first of these has 12 seats in the premium cabin and 138 in the economy cabin (for a total of 150 seats); the other option has 180 seats in a single economy cabin arrangement.

Aircraft Type	Configuration	First Class	Business Class	Main Cabin Extra	Main Cabin	Total
Boeing 767	Low Density	10	30	-	128	168
	High Density	-	28	25	156	209
Boeing 777-200	Low Density	16	37	-	194	247
	High Density	-	45	45	170	260

Figure 1-1: Major US airline Boeing 767 and Boeing 777 seats by aircraft

In order to forecast the operating revenue that will be generated by a specific aircraft type with a determined seat configuration, there are many assumptions that have to be made by the airlines. The first parameter to be forecasted is the traffic effectively carried by the airline on its flights. This traffic forecast (measured in Revenue Passenger Kilometers – RPK) would be affected by the future demand for travel between all the city pairs in the network, the expected market share of the airline in each of such markets, the projected load factors and by the available capacity (measured in Available Seat Kilometers – ASK) offered by the airline and its competitors. On top of the obvious effect of seat configuration of an aircraft on seat availability, it might also have an effect over other parameters mentioned above, such as the load factor. For example, an increase in the total capacity of an aircraft can possibly result in a higher share of travelers carried by the airline, but it does not necessarily imply that the traffic would increase in the same proportion as the increase in seat availability; i.e. load factor is likely to decrease. On the other hand, if the total capacity of an aircraft is reduced, potential demand is rejected and its revenue is not captured (this concept is known as spill and will be reviewed in detail in the following sections of this thesis). In dual cabin aircraft, an increase in the size of the premium cabin of an aircraft type will probably result in an increase in the share of the business passengers that will be off-set by a lower number of economy cabin passengers (and possibly of total passengers as well) transported by the airline.

The decision of how much capacity to allocate to each cabin might also have an effect on the other variable associated with the calculation of operational revenues: yield. Defined as the revenue per RPK, it is the indicator used for measuring the average price paid by the passengers for each kilometer flown. As explained in Section 1.1., premium cabins are installed on aircraft with the expectation of providing a product that incentivizes passengers with high willingness to pay to buy some of the most expensive fare classes available; therefore, the yield would be expected to increase because of the higher price paid by passengers for flying in premium cabin. However, the effect of it is not clear as adding premium capacity might also open space for lower-priced premium fare classes and reduce the availability for some of the highest economy fare classes; this situation is exacerbated by the fact that it is not unusual that the lowest restricted premium fare classes can have a lower price than an unrestricted economy fare class.

The implications of selecting a particular layout configuration also affect the costs of the airline. First, there is some effect on capital costs associated with the installation of different kind of seats in an aircraft type. Because of many factors, including but not limited to the installation of electronic mechanisms that allow premium seats to recline (even lay flat) and the use of customized materials, the unit price paid by an airline to a seat manufacturer for a premium cabin seat can significantly exceed the price of an economy cabin seat. Moreover, as airlines usually add special features to the seats in premium cabins (compartments for laptops, earphones, stowage, etc.), significant one-time fees are charged by the aircraft and seat manufacturer for the design, engineering, testing and certification of such features. In addition, the installation of better devices, such as larger (and more-expensive) in-flight entertainment monitors and handsets for an enhanced passenger experience can make the installation of premium cabin seats a very expensive proposition for an airline.

In addition, different configurations also have an impact on the operating costs of an airline; in that sense, the main effect is represented by the change in fuel consumption. As the physical characteristics of the seats are different, the unit weight of a premium class seat is significantly higher than the commonly simple economy cabin seats; however, since more economy cabin seats can be accommodated, the total weight would depend on the number of seats and their specific features. Of course, fuel consumption is also affected by the number of passengers effectively transported, as each one of them adds its own weight (plus its luggage) to the load carried by the aircraft. Compared to the total weight of an aircraft, these variations in weight can appear to be negligible, but it certainly represents a significant cost item to be considered when evaluating available configurations. Total service costs might also change as a result of the requirement of modifying the number of flight attendants working in the aircraft on each flight or in the provision of a different number of premium meals to be offered in each flight.

Operating costs and operating revenues are the main drivers for selecting an aircraft. Nevertheless, three other aspects involved in the seat configuration decision process are not so easily quantifiable or are more complex, but it is still relevant to mention them as part of the decision making process. First, as part of its brand positioning, an airline might consider strategically attractive to keep offering premium cabin services in order to project a world-class image and to be the preference of business travelers. The availability of space for upgrading passengers to premium cabin can be perceived as a mechanism that increases the fidelity of frequent flyers that usually buy high economy cabin fares. It could also be the case that a low cost carrier identifies the possibility of transporting business travelers if a premium cabin is incorporated in their aircraft, but the decision is to continue with a single cabin configuration in order to be consistent with the strategy of the airline. Second, in a similar approach to the previous item, an airline's decisions can also be affected by the cabin configuration of its competitors; just as with fares, airlines might need to "match" the services offered by another airline in order to remain competitive. Finally, although certain seat configurations might be favorable for most of the flights operated by an aircraft type, technical aspects such as the payload-range relationship can force the airline to operate the route with a different aircraft type or to restrict the amount of seats that can be made available for sale on such flights; as all the other aspects mentioned before, there could be a significant effect on the captured revenue as a result of this restriction.

Figure 1-2 summarizes the factors presented in this section as elements to be considered when analyzing available seat configurations for an aircraft type.

Revenue	Passengers by cabin carried on the fleet
	Passengers by cabin carried on the rest of the network
	Fares paid by the carried passengers
	Cabin capacity management
	Revenue Management
Costs	Fuel Consumption
	Configuration costs
	Other operational costs
Other variables	Brand and service considerations
	Comparison with competitors
	Aircraft technical characteristics

Figure 1-2: Factors considered in the evaluation of aircraft cabin configurations

1.3 Motivation for Research

As mentioned in the previous sections, the selection of an aircraft type in its fleet has long term effects on the profitability of the airline. In a market that keeps changing, such as the commercial aviation industry, any degree of flexibility for adjusting the capacity to match better the demand provides an opportunity for airlines to remain competitive when conditions have changed. If the market expands, airlines have the option of acquiring or leasing additional aircraft; if the demand for air transportation in the airline's market contracts, airlines might consider retiring or parking aircraft or early terminating leases.

Airlines can also explore alternatives that allow them to use more efficiently the "floor space" of their aircraft. In the short term, a carrier using at least two cabins in its configuration has the opportunity of sharing the capacity of the premium cabins, allowing for passengers that purchased economy fare classes to occupy the seats of such cabins when those seats are not expected to be filled by passengers paying high premium cabin fare classes. However, when capacity is shared there is a risk of displacing some of those passengers with high willingness to pay, reducing the revenue obtained from a sold seat. Another limitation of this approach is that it only allows the airline to increase the number of economy fare class passengers, but it does not provide an alternative for increasing the number of premium fare class passengers beyond the physical capacity of the premium cabin. A medium to long term solution that can be implemented by airlines consists of retrofitting aircraft in order to readjust the capacity of its cabins in order to match better the conditions of a market. The risks in this case are associated with a reduction of the total (premium + economy) revenue if the revenue lost from the cabin with reduced capacity is not compensated by the additional revenue gained in the enlarged cabin.

The evaluation of these alternatives would follow the same approach proposed in Section 1.2. Following such an approach, the strategy of sharing premium cabin capacity would probably have an effect over the airline's revenue with almost no effect on its costs. By contrast, the retrofit option would have parallel impacts on the revenue and the cost sides. Moreover, the implementation of a retrofit program has large additional costs involved, since aircraft would require to be grounded for the modification, aircraft LOPA requires to be certified by the aviation authorities and new seats and equipment have to be purchased for installation on the aircraft.

This thesis will focus on understanding the impacts that both strategies would have over the operating revenue of an airline implementing such changes. Instead of relying only on theoretical models for this evaluation, simulation tools will be used to address the complications related with the many assumptions required for calculating revenue (mostly the estimation of yield and traffic); more specifically, the Passenger Origin Destination Simulator PODS, described in Chapter 3 of this document, will be used. The advantage of this simulator is that, instead of considering that the changes only impact the airline implementing the strategy, PODS models a competitive environment where passengers make decisions based on their own characteristics and on the options made available by the airlines competing in the market; as a consequence, PODS allows us to evaluate the impacts on the decisions revenue of the airline and its competitors. Finally, it also presents the effect that a change implemented in a portion of the network (such as a specific aircraft type as part) has over the whole network of an airline, and allows for testing different revenue management techniques.

1.4 Thesis Outline

The remainder of the thesis is organized into several chapters: Chapter 2 gives a review of literature related with revenue management techniques, dual-cabin revenue management and spill modeling. Chapter 3 introduces the PODS simulation tool and presents the methodology that will be used in the thesis. Then, Chapter 4 explores the alternative of changing an aircraft's cabin configuration, and evaluates its impact on revenue. Chapter 5 is dedicated to the evaluation of the impacts of shared availability of premium cabin. Finally, Chapter 6 summarizes the results of the thesis, provides conclusions and suggests possible areas of future research.

2. Literature Review

This chapter summarizes the relevant literature and previous work related with the topics addressed in this thesis. First, a general framework of revenue management ("RM") is presented, focusing on the specific methods used in this study. Second, previous research on the multiple cabin revenue management problem is reviewed. Finally, literature related to spill modeling and its effect on the fleet assignment problem is introduced.

2.1 Revenue Management

As briefly described in Chapter 1, Revenue Management ("RM") systems are implemented with the objective of maximizing the revenue of an airline by capturing the highest possible revenue from each seat of each flight. As described by Belobaba et al. (2015), there are two main components of airline revenue maximization: differential pricing and revenue management. While the first offers different products at different prices in order to satisfy the demand of potential passengers with diverse willingness to pay, revenue management is the process used to managing the inventory of available seats between fare classes by setting booking limits on low-fare seats.

Demand forecasts for each fare class are one of the main inputs for RM systems. Such forecasts can be either for each flight leg operated by the airline or for each path itinerary. By combining these forecasts with the information of the prices associated with each fare class, RM optimizers and availability control mechanisms determine the expected revenue that would be captured if a seat in an aircraft is allocated to such fare class and/or path. The optimization process considers the trade-off between yields and loads so that the airline can control its seat inventory in a way that maximizes its revenue.

Leg-based RM uses information of the flight legs for its demand forecasts and its optimizer aims to maximize the revenue captured from each individual leg. The rules that are most frequently used in airline leg-based RM systems are based on the Expected Marginal Seat Revenue ("EMSR") model (Belobaba, 2016). Belobaba proposed the first version of this model, consisting of a heuristic that assumes that the demand for each fare class on each leg is independent of the demand for the other fare classes and that such demand is stochastic and follows a normal distribution (Belobaba, 1987). In addition, EMSR considers that all bookings arrive in a single booking period (static optimization) and that the lowest available class books before the next lower available class, and so on; it should be noted that the relevance of the last two considerations is reduced due to the possibility of calculating the booking limits many times during the booking period.

The EMSR model determines the number of seats to be protected for each fare class by comparing the marginal revenue that is expected to be captured by assigning one more seat to such fare class (considering its probability distribution) with the fare of each of the lower classes available; based on this algorithm, ESMR defines nested booking limits for each fare class. A modification to this heuristic proposes a generation of joint protection levels for higher fare classes relative to lower fare classes and is known as "EMSRb" (Belobaba, 1992). Although EMSR is a leg-based mechanism, it also controls connecting itineraries by restricting bookings in a certain fare class if such fare class is closed (i.e. does not have seats allocated as a result of the optimization process) on at least one of the flight legs included in the itinerary. As a result of this implementation, leg-based RM does not optimize the revenue captured from the network, since a passenger willing to book in a fare class for an itinerary that involves connections could be rejected even if the revenue that would be captured by accepting such a booking request is higher than the sum of the expected revenue to be captured on each of the legs.

This led to the development of network or Origin-Destination ("OD") revenue management; in contrast to the leg-based approach, OD RM focuses on maximizing the revenue of the whole network of the airline by considering local and connecting bookings simultaneously in the optimization. A key concept in this process is the "displacement cost" imposed by a connecting passenger by potentially displacing other passengers (and their revenues) from the legs included in the itinerary. Hence, the value added by a connecting passenger to the network is equal to the corresponding fare in the connecting market minus the displacement cost levied on the legs to be traversed (Belobaba et al., 2015).

One of the OD control mechanisms is known as Bid Price Control (Belobaba et al., 2015). This methodology determines a minimum price (i.e. the bid price) required to accept a booking request of a connecting passenger. The bid price for each leg traversed by a connecting itinerary is equal to the value of the last seat available on the leg plus the displacement cost on other legs that would be imposed by the connecting passenger on it. Moreover, in order to generate revenue gains for the airline, the fare of such a booking request must be higher than the sum of the bid prices of all the legs that are part of the itinerary.

The bid price control technique used in this thesis is Probabilistic Bid Price Control ("ProBP"). This method considers the probabilistic nature of the demand for each fare class and incorporates it by running EMSRb in each of the legs of the network, while considering the connecting itineraries that use the leg in the optimization as well. The value of such connecting itineraries is prorated between legs based on the marginal revenue values incremental available seats (or "EMSRc's") of those legs, and such EMSRc's are also used as the displacement costs for each leg. This process is repeated until the bid prices (displacement costs) converge for every leg in the network (Bratu, 1998).

2.2 Multiple Cabin Revenue Management

The multiple cabin revenue management problem aims to maximize the revenue of an aircraft with multiple cabins by sharing the seats of a higher service cabin with passengers who would book in the fare classes corresponding to a cabin that offers lower service standards. For example, in an aircraft with first class, business and economy cabins, first class seats could be shared with passengers that booked in business and economy cabin but business cabin seats could be shared only with passengers that booked in economy cabin. The main assumption for this approach is that passengers that booked in a certain cabin would always accept to be upgraded to a cabin that offers better service. Hence, the multiple cabin revenue management problem focuses on developing strategies to share the capacity of such higher quality cabins in a way that maximizes the revenue of the aircraft.

Alstrup et al. (1986) develop one of the earliest models found in the literature that splits the passengers in two types, "euro-class" and "tourist-class". Their model is a two-dimensional stochastic dynamic programming ("DP") that treats the booking process as a Markovian nonhomogeneous sequential decision process that includes overbooking, cancellations and no-shows. In addition, 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") that will be described in Chapter 3. Some of the evaluated strategies are based on a dynamic programming formulation and others on heuristics; as a matter of fact, these heuristics are closely related to part of the work that is presented in this thesis. The tested algorithms are the following:

- Multiple cabin DP Formulations:

    - Multiple Cabin DP:
      This approach assumes that the demand for each fare class in each time frame is independent and that booking requests arrive following a Poisson process. Bid prices are computed for each cabin using a DP algorithm. Upon arrival of a booking request for an economy fare class, its fare is compared with the bid prices. If the fare is lower than the economy cabin bid price but higher than the premium cabin bid price, the booking request is accepted and upgraded to premium cabin.
    - Multiple Cabin DP with variance:
      Similar to the Multiple Cabin DP, but this algorithm assumes higher variances by allowing arrivals of booking requests in batches and by not assuming equal variance for the different fare classes.

- Heuristics:

    - Shared nesting with EMSR:
        - Shared nesting Full EMSR:
          This heuristic applies the EMSRb algorithm to the entire capacity of the aircraft in order to determine the number of premium cabin seats to be protected for premium fare classes from economy fare classes. The booking limits for premium fare classes are adjusted based on premium cabin capacity while the booking limits for economy fare classes are left as calculated under the assumption of a single cabin.
        - Shared nesting Economy EMSR:
          In this case, the EMSRb algorithm is applied to the entire capacity and adjusts the booking limits for premium fare classes based on premium cabin capacity. The number of premium cabin seats to be shared (i.e. the premium cabin seats

that are not protected from economy fare classes) are assumed to be added to the economy cabin and EMSRb is applied again considering only the economy fare classes.

- Shared nesting EMSRc bid price control:
  Following the application of the EMSRb algorithm to each of the cabins separately, the EMSRc is calculated for each cabin and set as its applicable bid price. The number of premium cabin seats to be shared is determined by the largest number of premium cabin seats for which the economy cabin bid price is greater than or equal to the premium cabin bid price (i.e. the expected revenue of an additional seat in economy would be higher than expected revenue of the last available seat of premium). Once the number of premium cabin seats to be shared is determined, EMSRb is then applied adding the shared premium cabin seats to the economy cabin capacity.

- Multiple Cabin DP heuristic:
  Bid prices for each cabin are calculated using a DP formulation. In addition, an additional bid price is calculated for all the fare classes considering the total capacity of the aircraft.
  Premium fare classes booking requests are always compared to the premium cabin bid price. When the economy cabin has seats available, the economy fare class booking requests are compared with the economy cabin bid price; if economy cabin is sold out, the economy fare class booking requests are compared with the bid price calculated using the total capacity.

In order to compare the performance of the algorithms, these were tested by Lepage in two different networks and fare structures: on the one hand, a "single market" with two competing airlines operating one flight each; on the other hand, a "realistic" network with four airlines that compete in numerous flights and markets. The Multiple Cabin DP Formulations are not tested in this network because the author does not consider it practical due to the complexity and size of the scenario.

Based on the simulations, Lepage found that the Multiple Cabin DP heuristic leads to the highest revenue gains among the strategies evaluated; the revenue increase resulting from this approach can be as high as 2.5% (compared to distinct EMSR), but gains were observed in all the tested scenarios. In addition, Shared nesting Full EMSR reported revenue gains only in the "realistic" network (up to 0.4%). Revenue reductions were obtained from the simulations of all the other algorithms proposed. Overall, the heuristics outperformed the DP methods.

Walczak (2010) hypothesizes that higher flexibility in the mechanisms for sharing the premium cabin capacity results in higher revenue for the airlines. Such flexibility is provided either through availability of fixed size cabins, as in the multiple cabin revenue management problem, or through adjusting the aircraft configuration and the size of each cabin with a movable curtain. The author acknowledges that, in the case of the movable curtain approach, accepting one premium cabin passenger reduces the number of seats available to economy fare classes by more than one due to the seating arrangements.

For the problem in which the capacity of each cabin is fixed, two DP formulations are proposed. An insight from the solution of these formulations is that the output of the model is a two-dimensional matrix of bid prices that is not practical for inventory control because it cannot be implemented readily by existing systems. In addition, it is found that there is a relationship between the bid prices of one cabin and the available seats in the other cabin. 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 once economy is sold out.

Finally, 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. In a separate topic, the authors discuss how a pronounced and persistent mismatch between supply and demand that require the systematic use of upgrades may modify the expectations of the passengers, exacerbating the mentioned mismatch (the authors do not reach any conclusion in this respect).

2.3 Spill Estimation and Impacts on Fleet Assignment

As discussed in Chapter 1, traffic forecasts are a necessary input for assessing the expected revenues that would be captured by an airline if it decides to operate a certain aircraft configuration. Those forecasts are based on the estimation of the demand for air travel in each of the markets operated by the aircraft and on the expected market share captured by the airline in such markets. Observed traffic in the market can be used as one of the parameters for estimating the demand for air travel between two cities. In addition, Belobaba et al. (2015) identify socioeconomic and demographic variables (population, income per capita, etc.), trip purpose characteristics (mainly business versus leisure), prices of travel options (all the available modes of transportation) and other "quality of service" considerations (frequencies, waiting times, safety, etc.) as additional factors affecting the demand for air travel between two cities. Other models presented in the same book incorporate the concepts of price and time elasticity for demand for air travel in order to capture the sensitivity (or insensitivity) of different kinds of passengers to those variables.

It should be emphasized that in the context of passenger airlines there is a difference between "demand" and "traffic". Belobaba et al. (2015) conceptualize demand (at a given set of prices) as the sum of customers that are able to travel given the available capacity (i.e. effective "traffic") plus the consumers that were willing to travel but were not able to be accommodated in the available flights due to scarce seat capacity; this rejected demand is known as "spill". Following this approach, demand is always higher or equal than the observed traffic, so the associated spill is always higher (or equal) than 0. Considering the revenue losses resulting from rejecting a portion of the demand, the role of spill is of particular relevance when determining the adequate aircraft capacities in the aircraft fleet planning and assignment process.

One of the most commonly used methodologies for estimating spill is based on a model developed by Boeing (Boeing, 1978) for understanding the relationship between flight leg loads and passenger spill. Boeing develops a mathematical formulation that works as a base for deriving the demand distribution and estimating the spilled passengers on a flight leg (because of this, it is widely known as the Boeing Spill Model). The model assumes that the distribution of unconstrained total demand (i.e. ignoring the distribution between fare classes) for a flight leg follows a normal distribution; however, since the limited aircraft capacity allows the airline to observe only a portion of the total demand on flights when spill occurs, this normal distribution is truncated. A mathematical expression is provided to calculate the mean of the truncated normal distribution based on the observed load factors and an assumed coefficient of variation ("CV"); the CV describes the shape of the demand distribution and its typical value for the airline industry is suggested to be between

0.20 and 0.40. Methodologies based on the load information are also proposed to estimate the coefficient of variation.

Finally, Boeing provides normalized tables of spill as a function of load factors and demand factors for various values of CV. Based on its results, the Boeing Spill Model can be used to theoretically compare the passengers carried and the spill for any alternative capacity offered for each individual flight leg. Therefore, the recorded load distribution can be used by airlines when the aircraft capacity assigned for a flight leg is being decided, as part of the fleet assignment process. As presented by Subramanian et al. (1994), Delta considered the Boeing Spill Model as part of the Coldstart Project that addressed the problem of fleet scheduling.

One of the major shortcomings of the original Boeing Spill Model is that when it uses a single probability distribution for the demand to estimate the spill factor or spilled passengers, it ignores the role of fare classes and RM. Dealing with these shortcomings, Swan (1994) acknowledges that one effect of RM is that low-fare demand is turned away because of the protection of seats for high-fare demand; therefore, a flight does not have to have "full" capacity to reject passengers. In addition, the author references studies that suggest that the average spill fare should be a weighted combination of the lowest fare class (80%) and the average fare (20%). Abramovich (2013) references an unpublished paper by Swan that revisits the Boeing Spill Model and incorporates the effect of RM systems by aggregating the demands for the different booking classes, but still ignoring the booking limits computed by the RM system.

Li and Oum (2000) derive formulations for the expected spill and the spill rates based on four different probability distributions: normal, logistic, log-normal and gamma. The authors then proceed to compare the shape of the demand distribution when different CV values are used. Among other findings, it is reported that the tested distributions have similar behaviors for small CV values. On the other hand, when the CV value becomes larger, it is found that the normal distribution overestimates the true spill and that the difference in spill between the log-normal and gamma distributions is small. Other findings with high CV values are that the shape of the gamma is similar to the shape of an exponential distribution (not good for demand) and that neither a normal nor a logistic distribution are appropriate because these distributions would have a high probability of taking a negative demand (impossible under a Gamma or log-normal distributions). Using two numerical examples, the authors not only suggest that the choice of the demand distribution is a more important issue when demand is volatile but also that the capacity level is not so relevant in order to decide the distribution that should be used. Swan (2002) also contributes in this discussion by suggesting that Gamma shape might be relevant for first and premium cabins, where the mean values are much smaller than the capacities.

Based on insights provided in Farkas (1996), Farkas and Belobaba (1999) discuss the impact of OD RM on the estimation of spill on a single flight leg. As opposed to the aggregate demand approach of the BSM (and its Swan extension), the authors disaggregate the demand by fare class and by departure. The demand for each fare class is assumed to be independent of the demand of all the other fare classes; other assumptions include no cancellations, no recapture and that lower-valued classes book before the higher-valued fare classes. Combining this information with the booking limits previously provided by a RM optimizer, it is possible to estimate the spill by fare class; moreover, considering the fare values of each class the spill costs could be easily calculated instead of requiring estimating an "average spill fare" for the spilled passengers. An additional model that introduces a multi-period representation of fare class demand is also proposed.

Using Monte-carlo simulations of the booking process and under different discount ratios (i.e. the fare of a fare class over the fare of the immediately higher fare class), the proposed models are tested and compared with a method that uses an aggregated demand distribution. While the authors report that the differences between the two models based on disaggregated demand are negligible, the results suggest that spill is underestimated when aggregated demand is used as an input, potentially affecting the aircraft assignment process.

Abramovich (2013) examines some performance measures (total revenue, load factors, marginal revenue, bookings by fare class) resulting from different single cabin aircraft capacities. The scenarios simulate diverse passenger booking scenarios using PODS and implementing different RM systems, forecasting algorithms, fare structures and networks. The author initially evaluates a single flight network with no competition; in general, a decrease in load factors and in marginal revenue is observed as capacity is increased. Moreover, one of the main findings presented is that, under certain scenarios, increasing capacity can result in negative marginal revenues (i.e. more capacity available but less revenue captured by the airline). Similar results are observed when the airline operates two flights but there is still no competition and also when a second airline is introduced in order to evaluate the effect of competition in the results.

An additional set of simulations is presented in order to compare diverse spill models. Another heuristic, extending the Farkas spill model to incorporate sell up probabilities, is compared with the original Farkas spill model, the Swan extended Boeing Spill Model and the spill reported in the simulations. It is found that the proposed heuristic estimated the spill costs better than the Farkas and Swan-extended Boeing models, when compared with the simulated spills from PODS. In addition, the performance of the heuristic is better when a restricted fare structure is used than when a less restricted structure is in place.

2.4 Summary

This chapter has reviewed literature related with revenue management, the multiple cabin revenue management problem and spill modeling; the models and techniques presented herein provide a general picture of the context of this thesis. In addition, this review also helps to identify gaps in the literature that this thesis aims to fill. One of these gaps is identified after finding that there is no available literature that examines the effect of spill on aircraft with multiple cabins and how this information could be used to compare between configurations of an aircraft fleet operating in a network. In addition, this thesis also aims to contribute by proposing additional heuristics for the dual cabin multiple cabin revenue management problem. Finally, new insights can be provided to the revenue management literature by illustrating the impacts of using leg-based or OD-control RM in a dual cabin setting.

3. Methodology

This chapter introduces the methodologies that are proposed in this thesis and the tools used for evaluating the scenarios tested. Because of the RM insights provided by the Passenger Origin Destination Simulator ("PODS"), the first section of this chapter summarizes some of the main features of this simulator; this is accompanied by a brief characterization of the network that is used as the baseline for the results of this thesis. The last two sections are dedicated to introducing the proposed approaches for evaluating changes in seat configuration and shared capacity schemes, both of these in dual cabin aircraft.

3.1 Passenger Origin Destination Simulator (PODS)

Originally developed by Boeing, PODS is the simulation tool used for testing the outcome of the scenarios and strategies proposed in this thesis. Running on hypothetical airline networks, this software replicates the interactions between passengers willing to travel in diverse OD markets and airlines offering air transportation services in such markets. The fact that passengers in PODS are able to select between competitors is one of its most relevant features, as it allows us to evaluate the impact of the proposed strategies under the conditions of a competitive environment (like in the real-world).

The architecture of PODS consists of two major components: the Passenger Choice Model and the Revenue Management System. While the former models the passenger demand, characteristics and choices, the latter replicates the process from an airline's perspective, considering historical information, forecasts and RM optimizers. Interactions between these modules happen in both directions: on the one hand, the Passenger Choice Model feeds the airline's RM system when passengers book to fly in their chosen itineraries. On the other hand, the choices made by such passengers are based on the availability information provided by the airline, which is based on the Revenue Management System. The described structure is presented in Figure 3-1.

Another important characteristic of PODS is its consideration of the time dimension: the results of the simulation studies in this thesis correspond to the average result of 800 samples of a single departure date (details can be found on Tam et al., 2008). For each sample, 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, as presented in Table 3-1; as it is explained in the following sections, these time frames affect both components of the PODS structure.

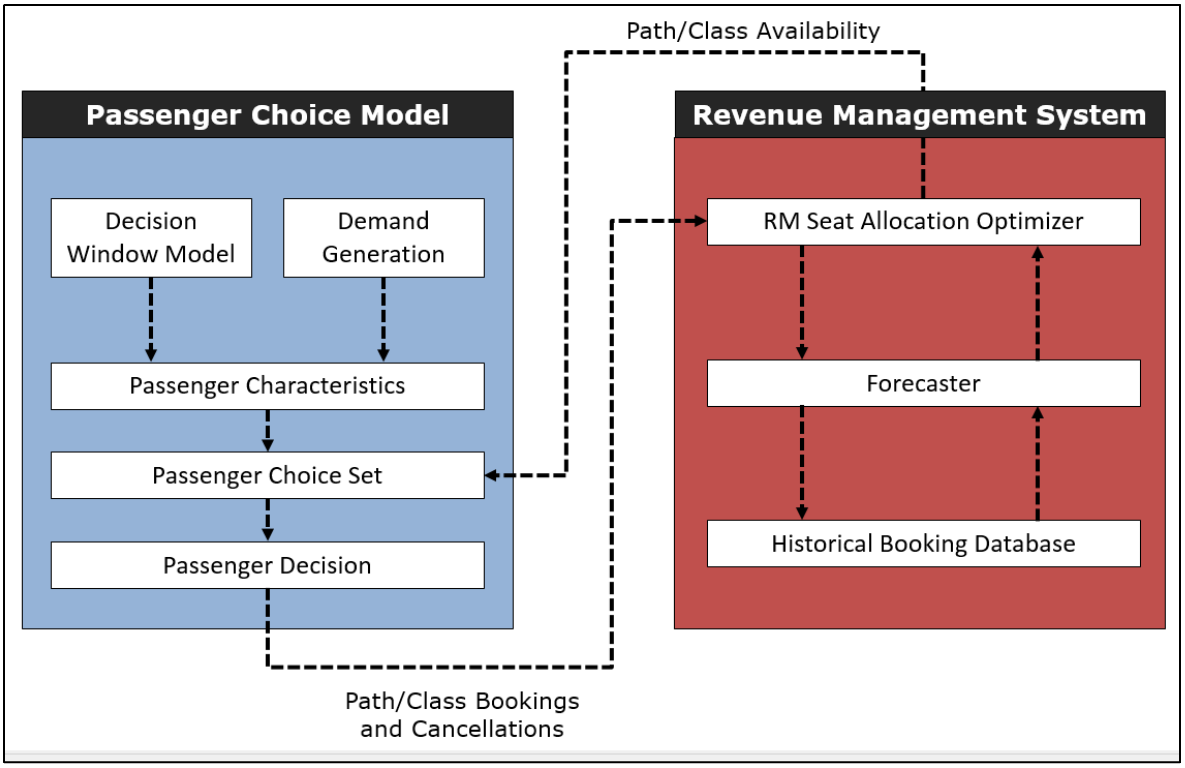

Figure 3-1: PODS Architecture

3.1.1 Passenger Choice Model ("PCM")

One of the sub-models of the PCM component of PODS is the model for generating demand. In this module, passengers are one of two types: business or leisure. For each passenger type in each market in the network, the demand for air transportation is generated considering an input demand and incorporating variability on a system, market and passenger type levels.

One of the main differences between these types of passengers is how their arrival process is modeled in PODS. On the one hand, leisure passengers are typically associated with their ability to make their travel plans well ahead of time and their interest in low fares, usually offered by airlines early in the booking period. On the other hand, very often business passengers have to make travel plans few days before their travel dates and are less concerned about paying high fares. Therefore, on average, a high proportion of the booking requests of leisure passengers is expected to be observed in the early time frames of the booking period, while most of the business passengers' booking requests arrive in the last stages of the booking period. These different arrival processes for each passenger types are modeled based on the curves presented in Figure 3-2, but variability in the arrival rate in each TF is also incorporated in PODS. It should be noted as well that although passengers are clearly identifiable as business or leisure in the PCM, the RM Systems of the airlines do not have the capability of identifying these passengers by type.

Time Frame	Days Until Departure	Time Frame Duration (days)
1	63	7
2	56	7
3	49	7
4	42	7
5	35	4
6	31	3
7	28	4
8	24	3
9	21	4
10	17	3
11	14	4
12	10	3
13	7	2
14	5	2
15	3	2
16	1	1

Table 3-1: Booking Process Time Frames

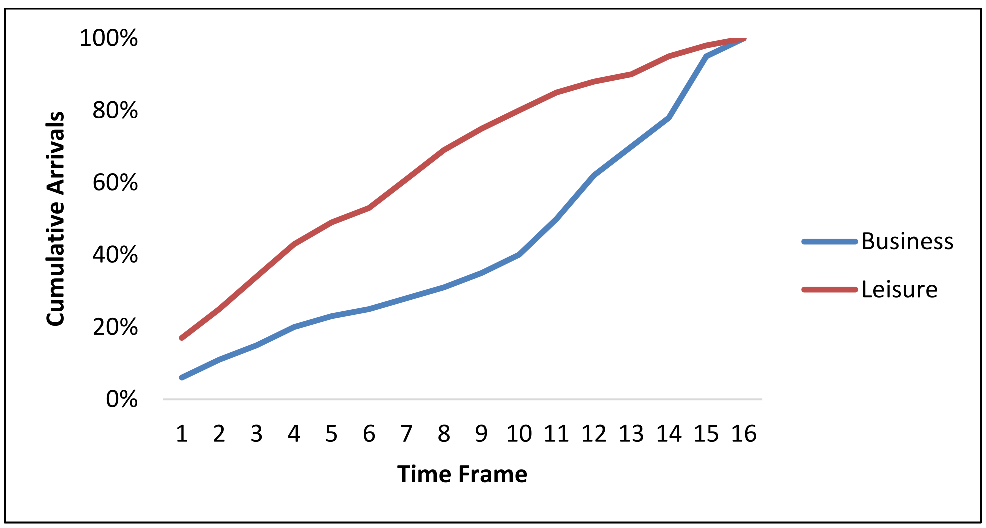

Figure 3-2: Demand Arrival by Passenger Type

The passenger type is also considered when assigning its unique characteristics to each generated passenger. The first of these dimensions of preferences is related to their willingness to pay

("WTP"); this value is generated for each passenger using a base fare and a probability distribution for the maximum proportional increase over such base fare that the passenger would be willing to pay for traveling. Differences between passenger types are then incorporated in the model by setting a higher base fare for business passengers than for economy passengers, resulting in higher average WTP values for the former.

Together with the WTP, each passenger is assigned a preferred time window for travel; hence, departing or arriving at a time that is not within such window generates re-planning costs for the passenger. The customer type is also relevant in this context, as re-planning costs are modeled to be higher (on average) for business travelers than for leisure passengers. The same idea applies for disutilities associated with travelling on an airline that is not the preferred or on an itinerary that requires connections (as opposed to flying on a non-stop flight); finally, each passenger has some level of inconvenience associated with the restrictions of each fare class. Although the specific value assigned to each passenger for each of these disutilities has a random component, the approach is the same: on average, an inconvenience caused to business passengers generates a higher level of disutility than a restriction imposed to economy passengers.

The choice set available for a passenger to travel from an origin to a destination is then reduced by two criteria: first, given the generated WTP of a passenger, the available options that have a fare that is higher than such WTP have to be discarded. In addition, airlines also can reduce the fare classes' availability as part of their revenue management practices. In summary, the choices available for the passenger are the ones that can be paid (lower than WTP) and that the airlines are willing to sell; it should be noted that not travelling is an option that is also available for passengers. If options are still available, the decision criterion that is incorporated in PODS is that passengers book in the option with the lowest generalized cost (i.e. sum of the fare plus the disutility costs indicated before).

3.1.2 Revenue Management System

The Revenue Management System ("RMS") is the component of PODS that simulates the airlines' process of deciding which air travel products are offered to their customers. As illustrated in Figure 3-1, the RMS consists of a historical database, a forecaster and an optimizer; this section introduces the forecasting and optimization mechanisms available in PODS that are used in this thesis.

- **Leg-based RM Optimizer: EMSRb**

A commonly used algorithm for calculating leg-based protection levels and booking limits is the Expected Marginal Seat Revenue ("EMSR") heuristic; as discussed in Chapter 2, this mechanism was developed by Belobaba in 1987 ("EMSRa") and then developed to "EMSRb" in 1992 (Belobaba, 2009). The objective of EMSRb is to maximize the revenue of an airline by defining joint protection levels for higher fare classes relative to lower fare classes. The EMSRb heuristic assumes that the demand for each fare class i is described by an independent Gaussian distribution; the mean ($\overline{X}_i$) and standard deviation ($\hat{\sigma}_i$) for such distribution are determined based on detruncated historical data. With these parameters, it is easy to calculate $P_i(S_i)$, the probability that the random demand for class i, X_i, is higher than the number of seats S_i made available to such class i. $EMSR_i(S_i)$ is also defined as the Expected Marginal Seat Revenue of making the Sth seat available to class i.

In order to calculate how many seats to protect jointly for classes 1 through n from fare class $n+1$, the following values have to be calculated:

$$\overline{X}_{1,n} = \sum_{I=1}^{N} \overline{X}_i$$

$$\hat{\sigma}_{1,n} = \sqrt{\sum_{i=1}^{n} \hat{\sigma}^2{}_i}$$

$$R_{1,n} = \frac{\sum_{i=1}^{n} R_i \times \overline{X}_i}{\overline{X}_{1,n}}$$

Then, it must be found the value of π_n that makes:

$$EMSR_{1,n}(\pi_n) = R_{1,n} \times P_{1,n}(\pi_n) = R_{n+1}$$

And the booking limit BL_{n+1} for fare class $n+1$ is set as:

$$BL_{n+1} = Capacity - \pi_n$$

The EMSR Booking Limit Optimization is illustrated in Figure 3-3: the booking limit for fare class B is set at the number of seats where the EMSR for class B equals the EMSR for class Y, meaning that if fare class B bookings were allowed beyond that limit, the expected revenue of each additional seat would be lower than when only fare class Y bookings are accepted. The figure also presents the concept of the "critical EMSR", $EMSR_c$, the EMSR value at the capacity of the leg; it also represents the expected revenue that would be lost if the capacity was reduced by one seat.

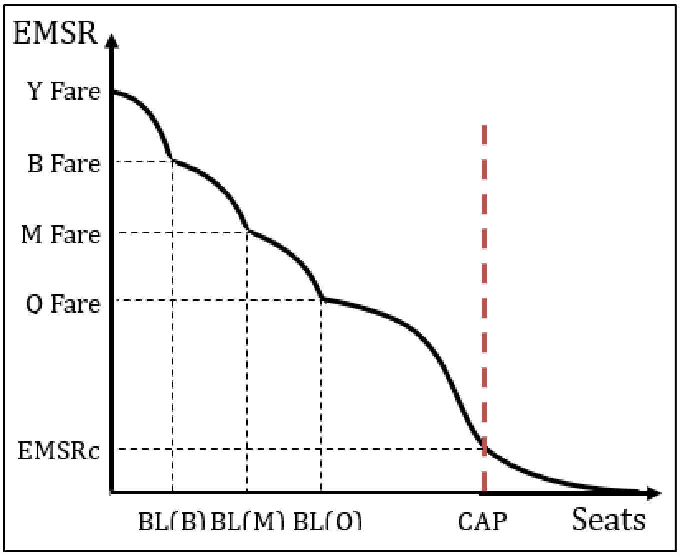

Figure 3-3: EMSR curve and booking limits

Being a leg-based RM method, EMSRb focuses on maximizing the revenue of a leg but ignores the effect of the determined protection levels over the network. However, this does not mean that the information of connecting passengers flying in the leg is not considered in the optimization; instead, all the demands for each fare class are aggregated (irrespective of their respective fare values) and used in the forecasting and booking level setting processes that are described in the following pages. In addition, connecting itineraries are controlled by restricting bookings on a certain fare class if such fare class is closed on at least one of the flight legs included in the itinerary.

- **OD-control RM Optimizer: Probabilistic Bid Price Control (ProBP)**

As briefly described in Chapter 2, Bid Price Control is one of the OD control mechanisms used by airlines to maximize the network revenue by considering local and connecting bookings in the optimization. ProBP is one of the mechanisms implemented in PODS to perform Bid Price Control and is the one that is used in this thesis. It incorporates the probabilistic nature of OD demand and requires as an input the fare and demand forecasts for each of the fare classes on each of the paths.

The ProBP algorithm for determining the bid price BP_k of each leg k is illustrated in Figure 3-4 (Bratu, 1998). Each iteration of the algorithm requires calculating $EMSRc_k$, the critical EMSR of leg

k and using this value as an input for determining the prorated fare $PRF_{j,k}$ of the origin destination fare ("ODF") j allocated to leg k. Defining L_j as the set of legs traversed by ODF j and F_j as the original fare of ODF j, such $PRF_{j,k}$ are calculated as follows:

$$PRF_{j,k} = \frac{EMSRc_k}{\sum_{m \in L_j} EMSRc_m} \times F_j$$

As shown in Figure 3-4, the ProBP algorithm iterates until the bid prices converge. The seat inventory is then controlled using the bid prices calculated for each leg k; booking requests for a fare class i in an itinerary j are only accepted if the corresponding fare $F_{i,j}$ is higher than the sum of the bid prices of the legs in set L_j.

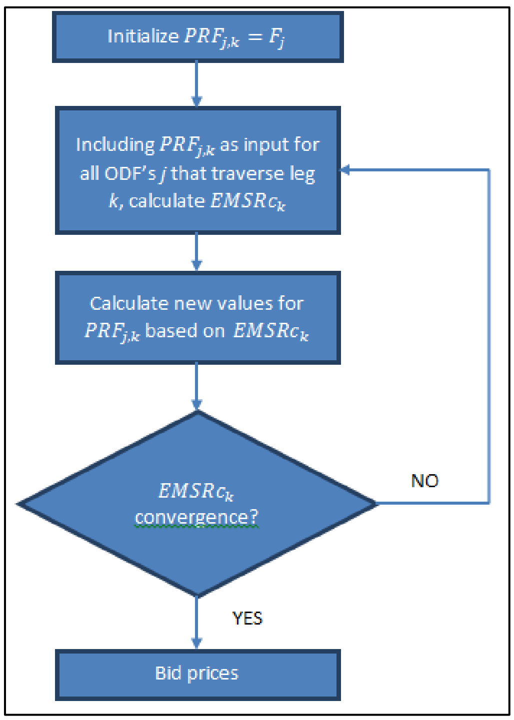

Figure 3-4: ProBP Algorithm

- **Forecasting**

Forecasts are the main input used by optimizers to calculate the booking limits or bid prices that are used on leg-based RM and OD-control RM, respectively. Among the forecasting methods available in PODS, two methods are used in the simulations of this thesis: standard leg-class forecasting and standard path-class forecasting. As suggested by its name, standard leg-class forecasting is the scheme used as an input for leg-control RM; in contrast, standard path-class forecasting is used for OD-control RM.

The main difference between leg-class and path-class forecasting is that while the former forecasts for each fare class on each leg based on the aggregated historical booking information available at the leg level, the latter forecasts for each fare class on each path using information available at the path level. However, the structure of both mechanisms is based on the pick-up forecasting methodology: it consists of forecasting, at each time frame, the total bookings at departure by adding the bookings in hand to the expected bookings to come in the remaining days of the booking period. The bookings to come are estimated based on the historical data available to the airline; as a matter of fact, the booking information observed after each sample in PODS is added to the historical booking database of the RMS of the airline, detruncated (in order to reflect demand constrained by booking limits) and considered in the forecasts used in the subsequent samples within the same trial.

Other forecasting mechanisms available in PODS are designed to support the RM process of the airlines by preventing the circuitous process known as “spiral down”, observed in unrestricted fare structures (Fry, 2015); despite the potential benefits of such alternate forecasting mechanisms, these are not incorporated in the simulations of this work since the insights are beyond the scope of this research.

3.2 Network V1

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 any of the 572 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 any of these destinations, either as the origin or as the destination, is catalogued as an international market; meanwhile, markets connecting two of the remaining 34 cities are considered domestic.

3.2.1 Airlines

All the airlines operate in a hub-and-spoke network structure, with a different hub for each; this structure is complemented by a handful of point-to-point legs offered by some of the airlines. The network of Airline 1 (“AL1”), Airline 2 (“AL2”), Airline 3 (“AL3”) and Airline 4 (“AL4”) are illustrated in Figure 3-5, Figure 3-6, Figure 3-7 and Figure 3-8, respectively. As can be observed in these figures, AL1 is the only airline that flies to all the cities in the network from its hub in MSP. AL2 has a similar route network (with its hub located in ORD), but it does not operate to every city in the network; however, it provides more point-to-point services in the network than AL1. AL4 operates 100% of its flights from or to its hub in DFW and also serves most of the cities. Finally, AL3 serves most of the domestic markets by flying to almost every domestic city, but it does not operate to any city beyond the boundaries of the contiguous United States of America.

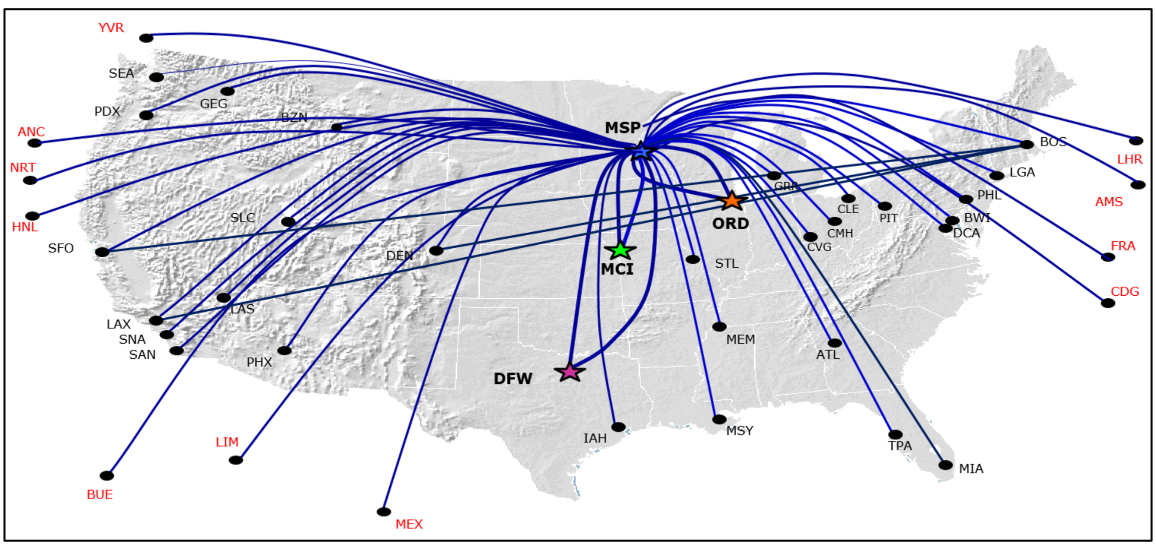

Figure 3-5: Airline 1 Route Map. Source: Belobaba, 2010.

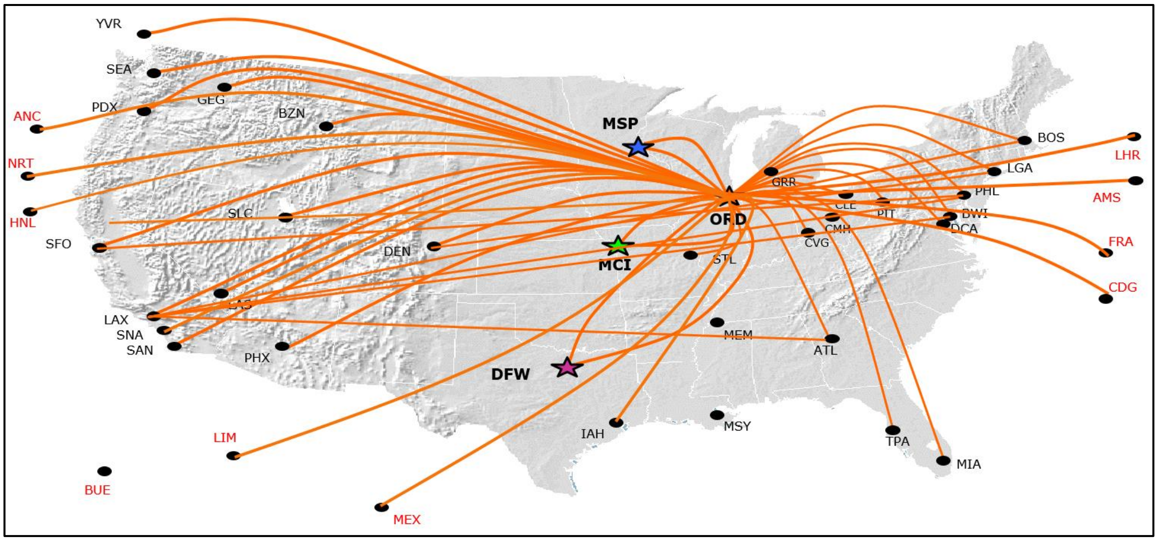

Figure 3-6: Airline 2 Route Map. Source: Belobaba, 2010.

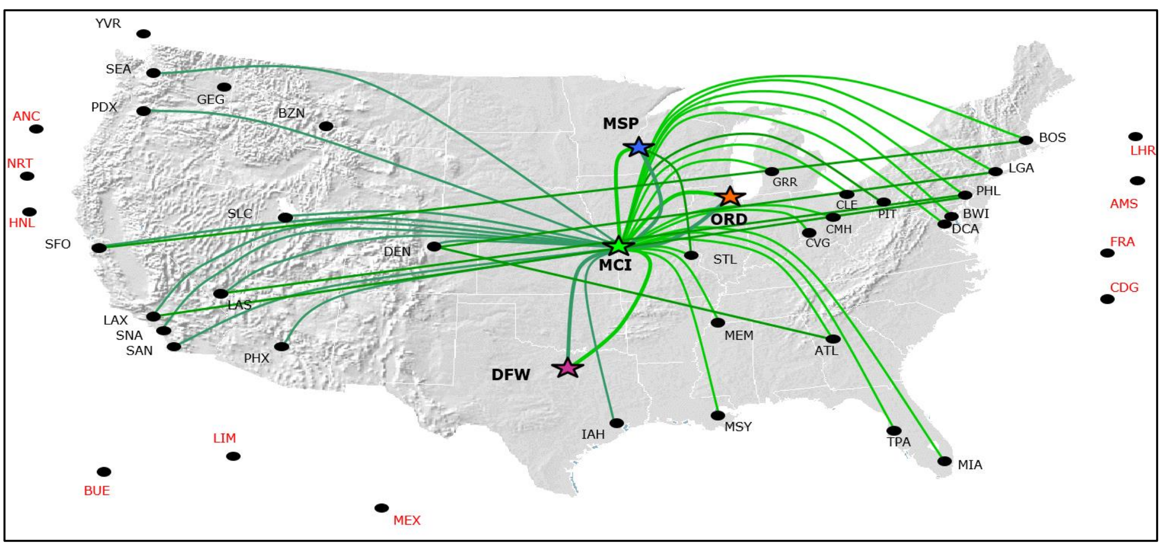

Figure 3-7: Airline 3 Route Map. Source: Belobaba, 2010.

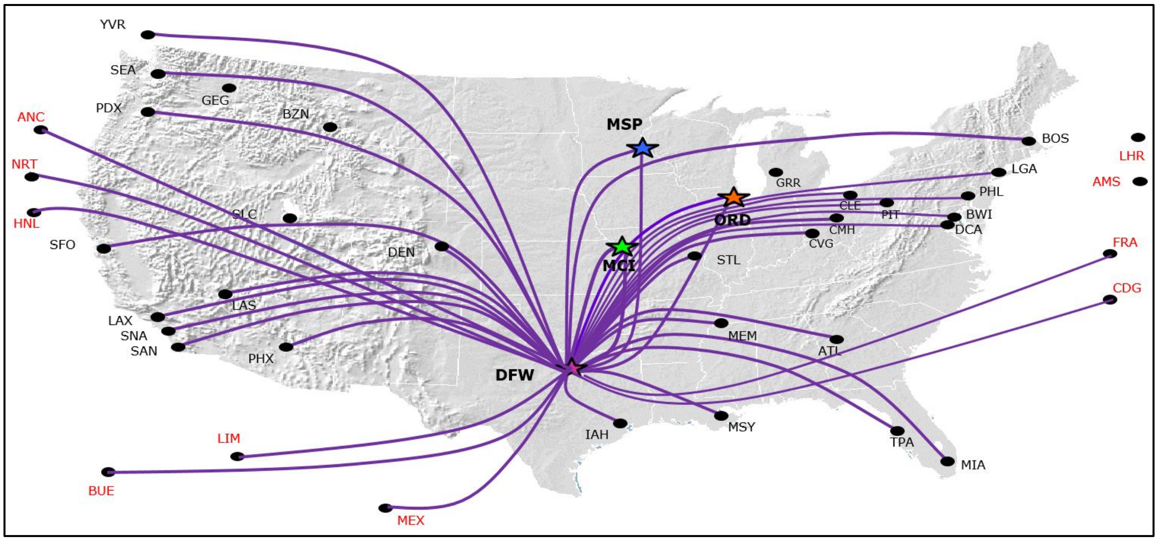

Figure 3-8: Airline 4 Route Map. Source: Belobaba, 2010.

The fleet composition and the operating profiles of each of the airlines are detailed in Table 3-2 and provide some important insights about the network. First, it shows that AL3 is the only airline that does not offer premium cabin seats; instead, it uses smaller aircraft that have the same economy cabin capacities than the economy cabins of certain aircraft types operated by the other airlines. Second, it illustrates the different operational profiles of the airlines; while AL1 and AL2 operate most of their flight legs in narrow body aircraft with capacities ranging from 112 (12 premium +100 economy) to 170 seats (20 premium + 150 economy), and complement it with a significant number of long distance flights operated in long haul aircraft with capacities exceeding 224 seats, AL3 operates most of its legs in short routes on aircraft with a capacity lower than 100 seats. AL4 operates an important number of flight legs both on short-haul aircraft with capacities up to 79 seats and on long haul aircraft, but most of its operation is on narrow body aircraft as well.

3.2.2 Fare Class Restrictions and Fare Structures

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 Network V1. Furthermore, these FCs are divided in two subsets: while FCs 1 to 4 give access to seats in premium cabin, FCs 5 to 10 are designated for passengers traveling in economy cabin (although this assumption can be relaxed, as proposed in the multiple cabin revenue management problem described in Chapter 2). As part of the differential pricing approach used by airlines in order to maximize their revenue, FCs with low fares are offered for passengers that can accept the applicable restrictions in exchange for affordable tickets; in contrast, FCs with high fares are typically designed for passengers that have high WTP but are less likely to accept restrictions.

Aircraft Size (Premium Capacity-Economy Capacity)	**1**		**2**		**3***		**4**	
	Legs	**Avg. Stage Length (mi)**	**Legs**	**Avg. Stage Length (mi)**	**Legs**	**Avg. Stage Length (mi)**	**Legs**	**Avg. Stage Length (mi)**
9-50	4	549	4	284	32	523	16	744
9-70	10	494	15	364	13	822	15	836
12-100	32	1,049	27	1,077	32	1,256	25	1,068
16-120	29	1,134	24	1,178	22	1,174	21	1,007
20-150	30	1,271	31	1,270	6	1,234	15	1,479
24-200	12	2,709	0	N/A	0	N/A	7	3,284
30-250	9	3,668	10	4,097	0	N/A	1	5,126
Total Legs/ Avg. Stage Length (mi)	**126**	**1,406**	**111**	**1,300**	**105**	**960**	**100**	**1,226**
*The Premium Cabin capacity for Airline 3 is 0 seats for all the aircraft types in its fleet								

Table 3-2: Fleet Composition and Operating Profile Airlines Network V1

As described previously in this chapter, the markets served by the airlines in network V1 are classified between two products: international and domestic. Other than the obvious geographical distinction between such products, there are also considerable differences between the fare restrictions and disutilities applicable to each of them. In order to explain these differences, it is necessary to introduce the FC restrictions and disutilities incorporated in PODS:

- AP: Advance Purchase requirement, ranging from 0 to 21 days.
- R1: Saturday night stay requirement
- R2: Cancellation or change penalty
- R3: Non-refundability
- R4: Disutility of Not Sitting in Premium Cabin
- R5: Disutility of Not Sitting in Premium Cabin on International Flights

Table 3-3 and Table 3-4 present the sets of conditions, restrictions and disutilities applicable for international and domestic products, respectively. Comparing between markets, it can be observed that international products have a more restricted fare structure than domestic products; another difference is that R5 applies only for international markets. Comparing within products, it is evidenced how 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 on 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.

	Fare Class	AP	R1	R2	R3	R4	R5
Premium	1	0	NO	NO	NO	NO	NO
	2	0	NO	NO	YES	NO	NO
	3	7	YES	YES	NO	NO	NO
	4	21	YES	YES	YES	NO	NO
Economy	5	0	NO	NO	NO	YES	YES
	6	3	NO	NO	YES	YES	YES
	7	7	NO	YES	YES	YES	YES
	8	10	YES	YES	NO	YES	YES
	9	14	YES	NO	YES	YES	YES
	10	21	YES	YES	YES	YES	YES

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

	Fare Class	AP	R1	R2	R3	R4	R5
Premium	1	0	NO	NO	NO	NO	NO
	2	0	NO	YES	NO	NO	NO
	3	7	NO	NO	YES	NO	NO
	4	21	NO	YES	YES	NO	NO
Economy	5	0	NO	NO	NO	YES	NO
	6	3	NO	YES	NO	YES	NO
	7	7	NO	NO	YES	YES	NO
	8	7	NO	YES	YES	YES	NO
	9	14	NO	YES	YES	YES	NO
	10	21	NO	YES	YES	YES	NO

Table 3-4: Network V1 - Applicable Conditions and Restrictions in Domestic Markets

Table 3-5 andTable 3-6 provide selected statistics that describe the fare structure for the international products; similarly, Table 3-7 and Table 3-8 offer information about the domestic products. As expected, fares are typically higher for international products than for domestic products, for both the premium and the economy FC's.

	FC1	FC2	FC3	FC4
Maximum Fare	$6,074	$4,555	$3,340	$2,429
Average Fare	$2,889	$2,167	$1,589	$1,156
Minimum Fare	$908	$681	$499	$363
Max. Fare Ratio	11.93	8.95	6.56	4.77
Avg. Fare Ratio	6.52	4.89	3.59	2.61
Min. Fare Ratio	3.43	2.57	1.89	1.37

Table 3-5: Network V1 - Fare Structure Premium Fare Classes in International Markets

	FC5	FC6	FC7	FC8	FC9	FC10
Maximum Fare	$3,037	$2,359	$1,670	$1,361	$1,180	$979
Average Fare	$1,144	$1,008	$799	$673	$574	$476
Minimum Fare	$454	$263	$178	$155	$133	$97
Max. Fare Ratio	5.97	4.28	2.70	2.09	1.43	1.00
Avg. Fare Ratio	3.262	2.27	1.77	1.45	1.22	1.00
Min. Fare Ratio	1.72	1.49	1.23	1.12	1.05	1.00

Table 3-6: Network V1 - Fare Structure Economy Fare Classes in International Markets

	FC1	FC2	FC3	FC4
Maximum Fare	$1,504	$1,353	$1,016	$663
Average Fare	$764	$671	$502	$373
Minimum Fare	$344	$300	$214	$147
Max. Fare Ratio	7.55	6.80	4.98	3.70
Avg. Fare Ratio	4.80	4.21	3.12	2.30
Min. Fare Ratio	2.57	2.24	1.81	1.46

Table 3-7: Network V1 - Fare Structure Premium Fare Classes in Domestic Markets

	FC5	FC6	FC7	FC8	FC9	FC10
Maximum Fare	$1,203	$865	$609	$451	$374	$337
Average Fare	$578	$405	$310	$244	$195	$161
Minimum Fare	$255	$153	$124	$96	$79	$64
Max. Fare Ratio	6.04	4.53	3.09	2.26	1.43	1.00
Avg. Fare Ratio	3.62	2.50	1.91	1.51	1.21	1.00
Min. Fare Ratio	1.91	1.41	1.14	1.12	1.04	1.00

Table 3-8: Network V1 - Fare Structure Economy Fare Classes in Domestic Markets

RMS used by airlines in Network V1 record, forecast and optimize based in these FC's; however, it should be clarified that Airline 3 does not offer FC's 1 to 4 as it does not operate aircraft with premium cabin.

3.3 Baseline Scenarios

Some scenarios must be set as references in order to evaluate the performance of the mechanisms proposed in this thesis. Because PODS considers competition between airlines in the simulation, any decision taken by one of the airlines would affect the results obtained by all the other companies. As a result, the effects of a decision taken by one airline in PODS can only be measured reliably if all the other airlines keep making their decisions as in the base case; if not, the interaction between the decisions of more than one airline would complicate the analysis of the performance of any of such decisions.

Because of the characteristics of its network, all the experiments in this thesis consist on shifting some of the conditions simulated for AL1; hence, the network and revenue management systems for AL2, AL3 and AL4 remain unchanged in the simulations. In addition, since the effects of the decisions made by AL1 are expected to depend on the seat inventory control mechanism used in its RMS, different baselines are set for evaluating scenarios depending if it is uses leg-based or OD-control methods. The main characteristics of the baseline scenarios are summarized in Table 3-9.

Leg-based Base Case	OD-Control Base Case
Network V1	**Network V1**
AL1: Distinct EMSRb	**AL1: Distinct ProBP**
AL2: Distinct ProBP	AL2: Distinct ProBP
AL3: Distinct EMSRb	AL3: Distinct EMSRb
AL4: Distinct ProBP	AL4: Distinct ProBP
Leg-based Std. Forecasting	Path-based Std. Forecasting

Table 3-9: Baseline Summary

It should be noted that all airlines are performing distinct RM for their aircraft cabins in the baseline. The implication of this is that RM considers each of the aircraft cabins (premium and economy) independently of the other, as if each cabin was an independent aircraft. 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). An alternative to this approach consists of sharing all or part of the capacity of the aircraft; the implications of such shared capacity strategies are discussed in Section 3.5 and in Chapter 5.

In addition to setting the different RM systems for all the airlines in the network, a demand level dimension is introduced in the baseline in order to evaluate the effects of any of the proposed changes under different market demand conditions. Following this approach, three levels of demand –low, medium and high-- are also considered when defining the baselines. Summarizing, for each control mechanism (leg-based control and od-control) three demand levels are defined, for a total of six baselines; the main statistics of each of these baselines are presented in Table 3-10 and Table 3-11. In order to illustrate the relative importance of each cabin-product combination for each airline, Figure 3-9 and Figure 3-10 respectively present the share in passengers and revenue of each of these categories.

Demand Level	AL	RPM	Premium Revenue	PC LF%	Economy Revenue	EC LF%	Total Revenue	System %LF	Yield	Market Share %
Low	1	23,496,212	$684,377	53.41	$3,108,571	76.54	$3,792,944	73.97	0.1614	32.61
	2	19,601,293	$544,821	52.57	$2,452,625	79.02	$2,997,443	76.07	0.1529	28.47
	3	7,244,415	$0	N/A	$986,949	73.81	$986,949	73.81	0.1362	17.93
	4	13,135,599	$340,043	47.62	$1,798,119	76.58	$2,138,159	73.3	0.1628	20.99
Medium	1	25,274,214	$753,993	57.68	$3,368,368	82.3	$4,122,359	79.57	0.1631	32.63
	2	20,900,865	$595,077	56.42	$2,649,812	84.22	$3,244,888	81.11	0.1553	28.26
	3	7,919,327	$0	N/A	$1,066,058	80.69	$1,066,058	80.69	0.1346	18.09
	4	14,170,268	$383,699	52.52	$1,954,981	82.47	$2,338,680	79.08	0.165	21.02
High	1	26,996,575	$829,675	62.72	$3,617,220	87.77	$4,446,902	84.99	0.1647	32.68
	2	22,101,256	$654,355	61.33	$2,874,625	88.84	$3,528,983	85.77	0.1597	28.31
	3	8,566,669	$0	N/A	$1,149,083	87.29	$1,149,083	87.29	0.1341	18.21
	4	14,867,692	$425,973	57.84	$2,124,949	86.18	$2,550,922	82.97	0.1716	20.81

Table 3-10: Base Case Outputs with AL1 using EMSRb

Demand Level	AL	RPM	Premium Revenue	PC LF%	Economy Revenue	EC LF%	Total Revenue	System %LF	Yield	Market Share %
Low	1	23,528,986	$703,797	53.28	$3,105,724	76.67	$3,809,521	74.07	0.1619	32.48
	2	19,560,957	$544,708	52.58	$2,448,467	78.84	$2,993,175	75.91	0.153	28.45
	3	7,308,432	$0	N/A	$988,994	74.47	$988,994	74.47	0.1353	18.07
	4	13,120,650	$340,330	47.81	$1,795,965	76.46	$2,136,295	73.22	0.1628	21
Medium	1	25,246,589	$774,366	57.54	$3,366,013	82.22	$4,140,379	79.48	0.164	32.27
	2	20,884,452	$596,414	56.6	$2,643,641	84.12	$3,240,055	81.05	0.1551	28.35
	3	8,117,388	$0	N/A	$1,067,966	82.71	$1,067,966	82.71	0.1316	18.35
	4	14,120,051	$384,209	52.82	$1,953,669	82.11	$2,337,878	78.8	0.1656	21.03
High	1	26,835,031	$845,285	62.05	$3,639,989	87.28	$4,485,274	84.48	0.1671	32.33
	2	22,055,465	$652,392	61.28	$2,856,508	88.65	$3,508,900	85.59	0.1591	28.41
	3	8,725,636	$0	N/A	$1,146,562	88.91	$1,146,562	88.91	0.1314	18.35
	4	14,825,233	$424,669	57.86	$2,118,938	85.91	$2,543,607	82.73	0.1716	20.91

Table 3-11: Base Case Outputs with AL1 using ProBP

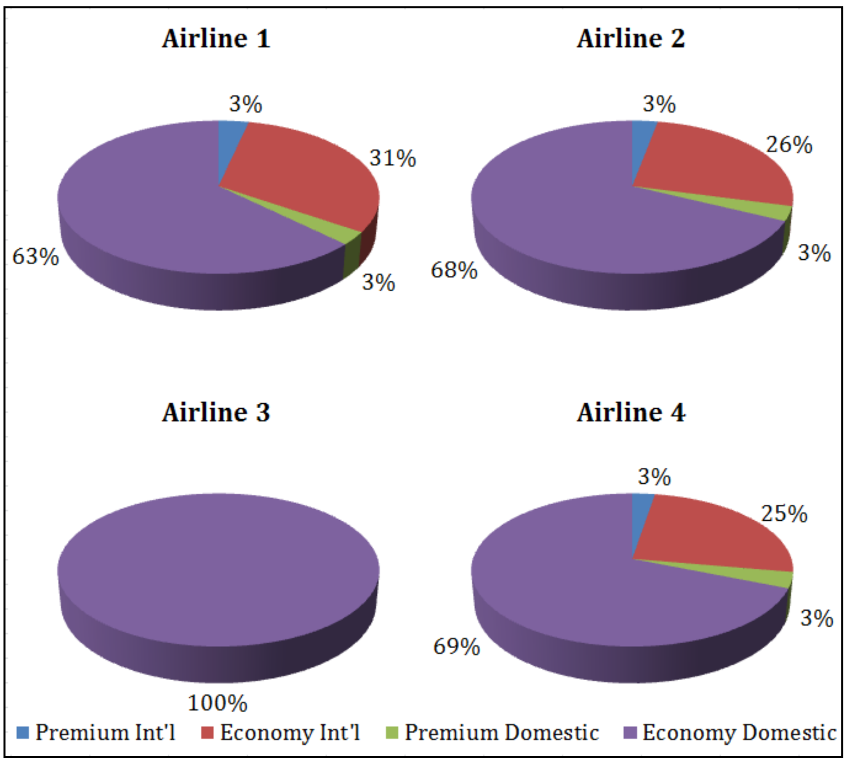

Figure 3-9: Airline passenger share by cabin and product – Leg-based RM Medium Demand Baseline

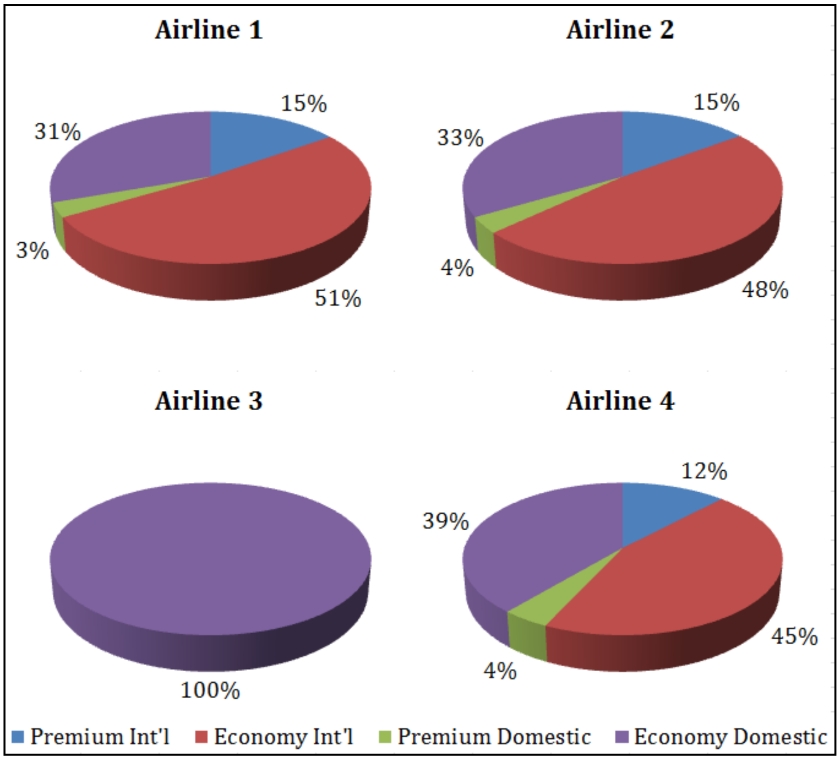

Figure 3-10: Airline revenue share by cabin and product– Leg RM Control Medium Demand Baseline

3.4 Cabin Configuration Analysis

As discussed in Chapter 1, fleet planning decisions related to the incorporation of an aircraft type in an airline's fleet are typically made under conditions of high uncertainty with regard to the market conditions that will prevail at the time when such fleet will be operated. Furthermore, these decisions affect the airline's performance for many years since the incorporation of an aircraft type is usually made for at least 10 years of operation. Airlines can adjust the capacity offered in a market by adding or retiring aircraft to/from their fleet; however, this process can be slow, expensive and complicated, as it involves negotiating with aircraft manufacturers or lessors and potentially implicates long lead times, or the payment of significant additional fees or penalties.

Subject to the technical characteristics of an aircraft type, the definition of its layout of passenger accommodation (LOPA) provides an additional degree of flexibility for airlines. Either by defining it just two years before receiving an aircraft from a manufacturer or by undertaking retrofits when the fleet is already in operation, airlines can modify the capacity made available in their markets to adjust better to the observed demand. Although changing the seat configuration of an aircraft type affects the costs of the airline and could potentially have implications on the position of the brand in the market, this thesis concentrates only on the impacts of such changes on an airline's revenue.

Intuitively, adding capacity to a single cabin aircraft should result in an increase in the revenue captured from each flight operated by such aircraft; likewise, it is expected that revenues should decrease if the number of seats is reduced. It could be also said that more demand could be accommodated in an aircraft with more seats, or that more spill should be observed if the capacity is reduced: the spill models introduced in Chapter 2 aim to provide theoretical approaches to the measurement of such spill. However, the impact of factors such as competition, sell-up and recapture can have a significant effect on the revenue obtained from a leg, as found by Abramovich (2013).

3.4.1 Seat Configuration Changes on Dual Cabin Aircraft

The cabin configuration analysis has an additional dimension of complexity on dual cabin aircraft: adding seats in premium cabin reduces the seats in economy cabin (and vice versa); thus, reducing spill in one cabin increases spill in the other. Moreover, because of the larger size of premium cabin seats compared to economy cabin seats and the increased pitch offered in the premium cabin, one seat of premium cabin capacity is exchanged by more than one economy cabin seat. Figure 3-11 and Figure 3-12 illustrate how adding premium cabin seats reduce the total capacity of the aircraft: while in a hypothetical narrow body aircraft the addition of a row with four premium cabin seats reduces the total capacity of the aircraft by eight seats, sixteen additional economy cabin seats can be accommodated by removing one row of six premium cabin seats in the case of a proposed wide body aircraft.

In this thesis, this topic is explored by modifying the seat configuration of the whole fleet of three different aircraft types operated by AL1 in Network V1 in PODS, as proposed in Table 3-12. Each of these fleet modifications are tested independently (i.e. only one aircraft type is changed in each simulation) and compared to the baselines presented in section 3.3. With the intention of understanding the impact of a configuration change not only on the legs operated by the modified aircraft but also in the rest of the network, each configuration is tested using leg-based control (distinct EMSRb) and OD-control (distinct ProBP); in addition, three different levels of demand (low, medium and high) are considered.

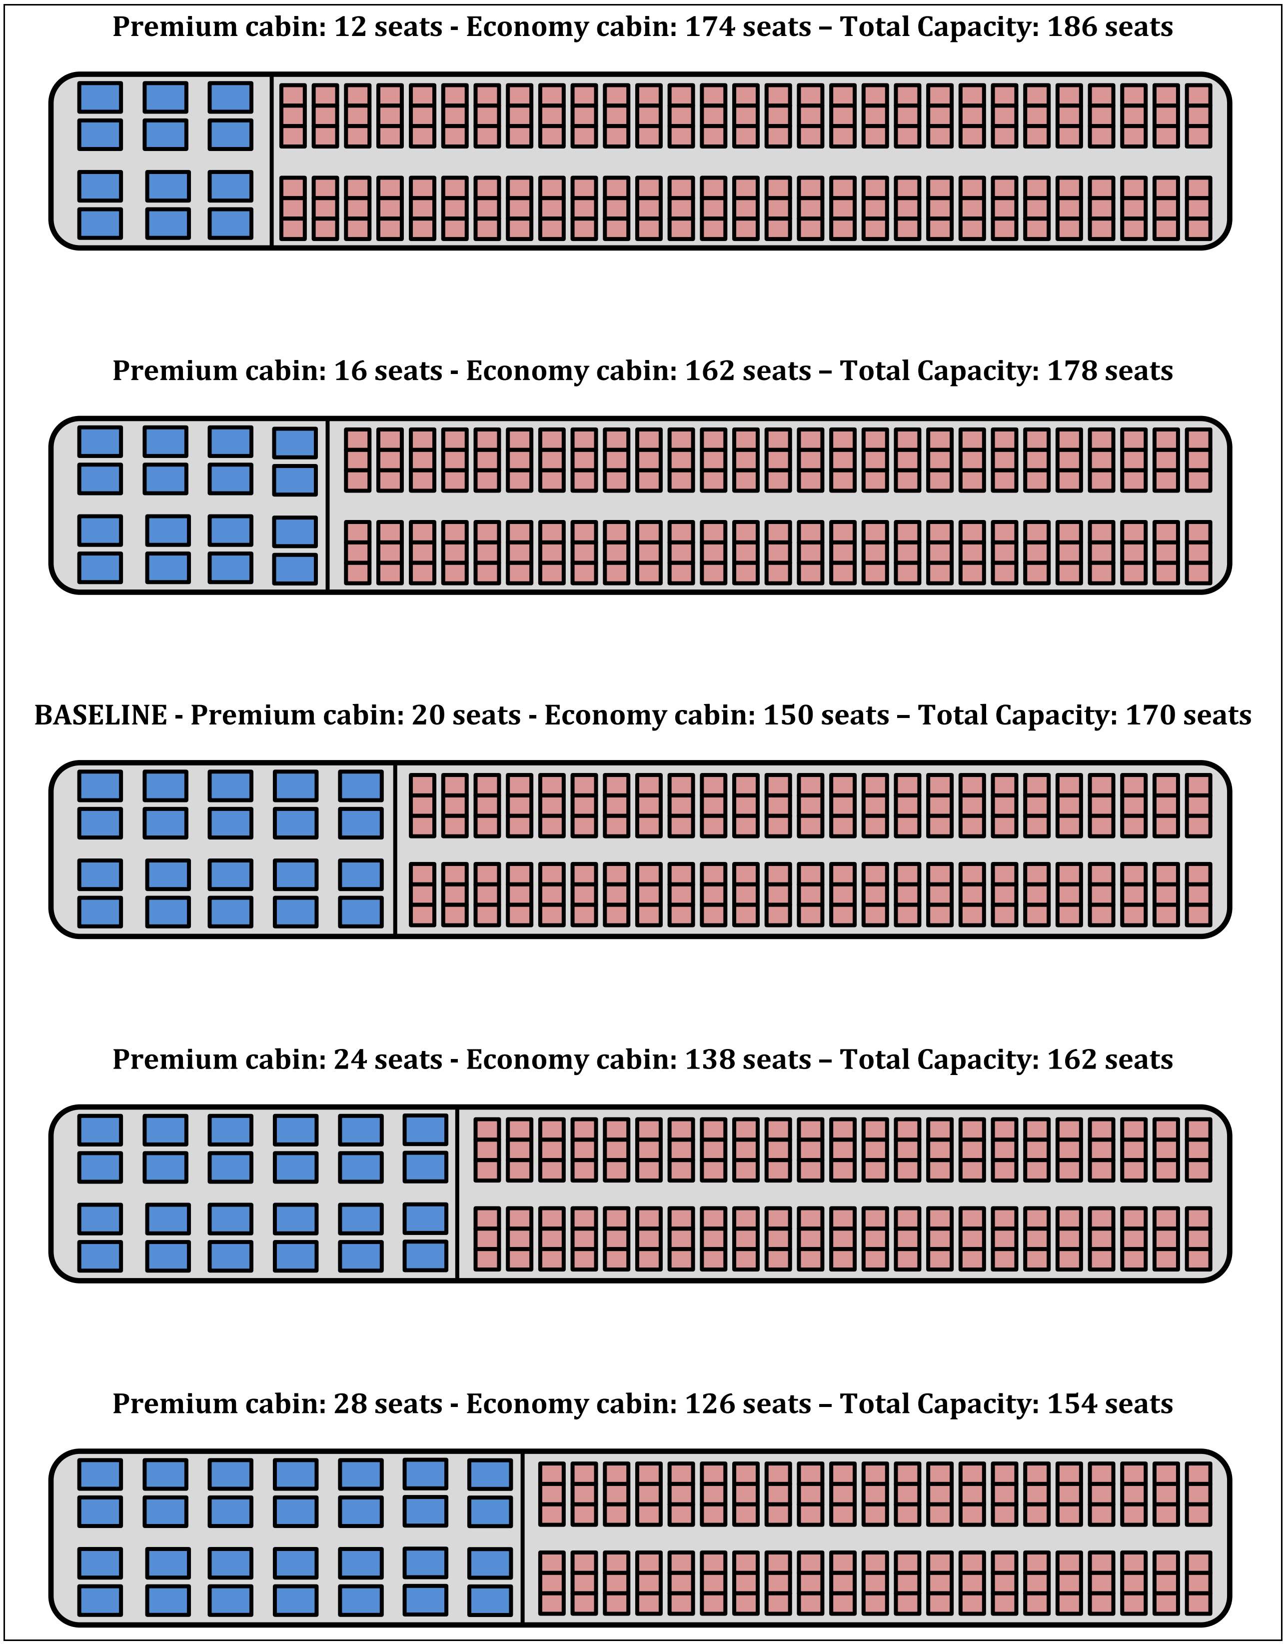

Figure 3-11: Alternate configurations for baseline configuration PC 20 – EC 150

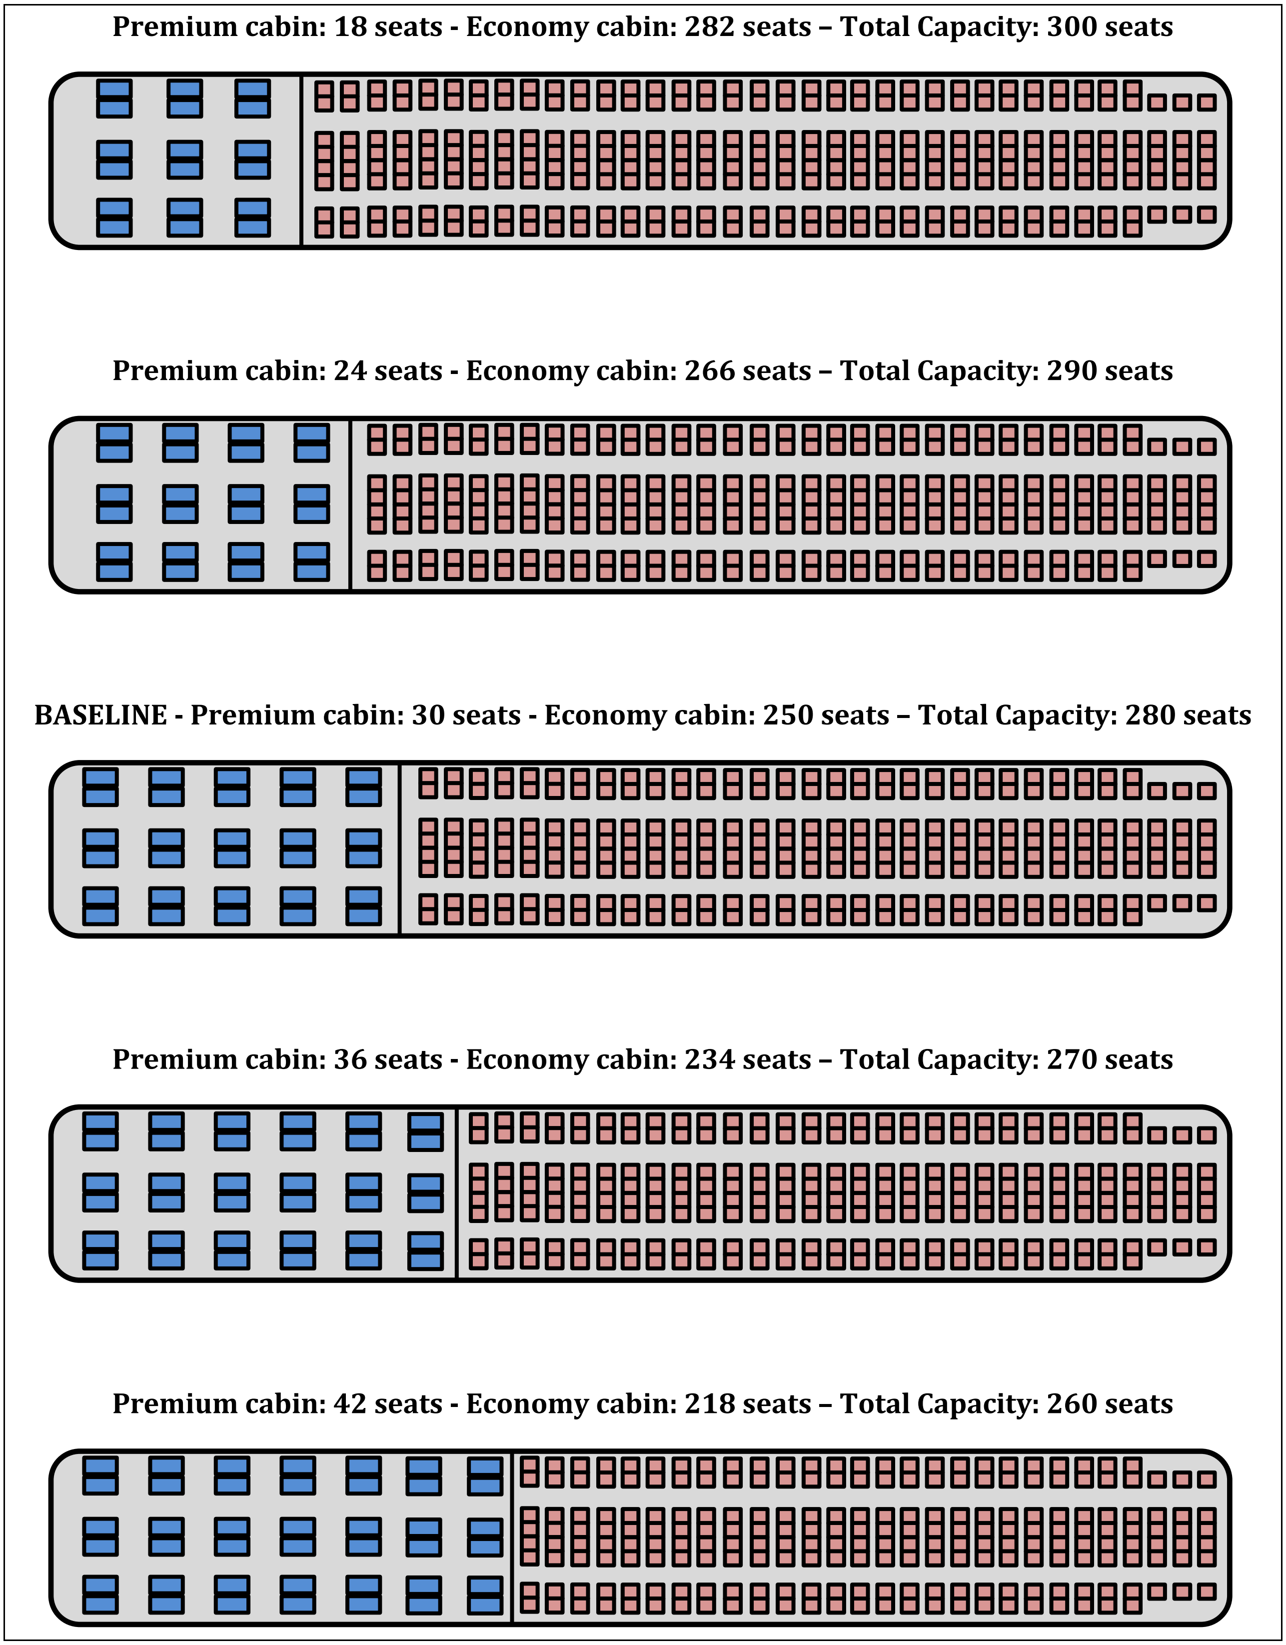

Figure 3-12: Alternate configurations for baseline configuration PC 30 – EC 250

A/C Type	Configuration High Premium (HP)	Configuration More Premium (MP)	Configuration Baseline (BL)	Configuration More Economy (ME)	Configuration High Economy (HE)
A	12-174	16-162	**20-150**	24-138	28-126
B	12-232	18-216	**24-200**	30-184	26-168
C	18-282	24-256	**30-250**	36-234	42-218

Table 3-12: Aircraft Seat Configurations Simulated

3.4.2 Spill analysis

One of the motivations of this thesis is comparing the results obtained in the simulations with the results that would be obtained by an airline when analyzing the potential impacts of changing the seat configuration of an aircraft fleet. As discussed by Belobaba et al. (2015), the magnitude of the spill can be estimated from observed loads using spill models; in this work, the model used is the Boeing-Swan Spill Model, described in the following section.

Other spill models found in the literature (and briefly discussed in Chapter 2) incorporate RM considerations such as fare classes and booking limits. However, most of those models require knowledge of the probability distributions for the demand at the fare class level but do not provide an analytical method for determining such distributions. Estimates of unconstrained demand by fare class could be made by considering the first choice demand for every passenger that is provided as an output of the simulations in PODS but this information is not available for the airlines in the "real-world" and therefore is considered beyond the scope of this thesis.

- **Boeing-Swan Spill Model ("BSM")**

The BSM assumes a single normally distributed function $f(x)$ with mean μ and standard deviation σ for the unconstrained demand of each leg. Conceptually, spill occurs when the demand is larger than the capacity of the aircraft: therefore, the expected number of spilled passengers can be expressed mathematically as follows:

$$SP = \int_{C}^{\infty} f(x)(x - C)dx$$

where SP is the total number of spilled passengers and C is the capacity of the aircraft.

It should be highlighted that $f(x)$ is the function for unconstrained demand; however, since all the observations of loads captured by the airline are constrained to some degree by the aircraft capacity (the constraints imposed by the booking limits in RM systems are not considered in the initial version of this model), airlines do not have complete information available for such unconstrained demand.

In order to estimate the mean of the unconstrained demand from the observed loads, Boeing (1978) provides an expression that uses the observed load factors and an estimation of the coefficient of variation of the unconstrained demand as parameters for calculating the mean demand. The expression is:

$$L = (D - 1) \times F_o\left(\frac{1}{K \times D} - \frac{1}{K}\right) - K \times D \times f_0\left(\frac{1}{K \times D} - \frac{1}{K}\right) + 1$$

Where

$$K = \frac{\sigma}{\mu}$$

$$f_0(x) = \frac{exp(-x^2/2)}{\sqrt{2\pi}}$$

$$F_0(x) = \int_{-\infty}^{x} f_0(t)dt$$

$$L = \frac{Average\ Load}{Capacity} = Load\ Factor$$

$$D = \frac{Average\ Demand}{Capacity} = Demand\ Factor = \frac{\mu}{C}$$

Based on this formulation and assuming a coefficient of variation K, the spill factor S and the expected spill SP are:

$$S = D - L$$

$$SP = C \times (D - L)$$

In order to simplify the calculation of D and S from observed L, Boeing provides tables for various levels of K. Based on the recommendation of using a factor K between 0.2 and 0.4, in this thesis the factor K used for both cabins is 0.3.

The reason why the BSM is of particular relevance to the topic of this thesis is that, in addition to calculating spill, this model can be easily used to estimate the passengers that would be carried if the capacity of the aircraft operating the flight was modified. Assuming a factor K for a flight operated with a capacity C_1 in which a historical load factor L_1 has been observed, a demand factor D_1 can be obtained from the model above (or from Boeing's tables), and therefore an estimation of the mean of the unconstrained demand, μ. By modifying the capacity to be C_2, a new demand factor D_2 can be calculated; L_2 can be determined from the tables, indicating the expected loads $L_2 \times C_2$ to be carried with the new capacity.

Although it is relevant to estimate the number of additional passengers carried if the capacity of an aircraft serving a route is increased, what really matters for airlines (as profit maximizers) is determining the impact on revenue resulting from such change. Therefore, it is also important to have an estimate of how much would be paid by those additional passengers. Swan (1994) suggests that, if RM systems are doing their job, most of the spilled demand would pay the minimum fare in the market; following this approach, the proposed fare of a spilled customer is a weighted average of 80% the lowest fare in the market and 20% the average fare in the market. Hence, this weighted average is used in this thesis as the value of each passenger accommodated or spilled, as calculated using the BSM.

3.5 Premium Cabin Capacity Sharing Mechanisms

As discussed in previous sections, dividing the aircraft in cabins with different service standards helps airlines to be more effective in their objective of providing different products to the broad range of values of willingness to pay of air transportation consumers. Nevertheless, such division also complicates the revenue maximization problem of the airline, as it imposes a restriction on how the aircraft seats can be allocated between passengers. An ideal scenario would be that in which the airlines are able to effectively maximize the revenue of the aircraft by maximizing the revenue of each of the cabins but only carrying passengers in the cabin corresponding to the fare class of their booking. However, the reality is that on many flights there is demand for economy fare classes that is spilled because seats in economy cabin are unavailable, although premium cabin seats are still available. The same situation can happen in the other direction, when passengers that would fly only in premium cabin are not able to be accommodated in such cabin while the economy cabin still has seats to offer.

In order to reduce the impact on revenues of the constraint imposed by the division of the aircraft into sub-cabins, sharing schemes can be proposed for taking advantage of available premium cabin capacity when the expected demand for economy cabin seats is higher than the number of available economy cabin seats. The premium capacity sharing methodologies explored in this thesis work under the assumption that passengers who booked in economy fare classes would be pleased to enjoy the higher service quality offered in premium cabin. In order to prevent misinterpretations we need to make two clarifications: first, economy cabin seats are not shared with passengers that booked in premium fare classes since these passengers would not accept a product of a lower quality for the price they have paid. Second, the upgrades discussed in this thesis 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.

Given the structural differences between Leg-control RM and OD-Control RM, the capacity sharing schemes that are presented in the following sections of this chapter are adjusted to the characteristics of each of such RM control mechanisms.

3.5.1 Capacity Sharing with Leg-Control RM

As specified in section 3.3, the baseline implementation of leg-control EMSRb applies this methodology distinctly in each of the cabins, as if these were different aircraft. Therefore, while capacity in the premium cabin is protected for the high fare classes only from lower premium fare classes, economy cabin capacity is protected from the low economy fare classes only for the higher economy fare classes. Consequently, a modification to this approach is required to share the capacity of premium cabin.

The proposed leg-control RM premium cabin capacity sharing approach is based on the Shared Nesting "Full EMSR" methodology proposed by Lepage (2013); basically, it consists of applying EMSRb to the entire capacity and modifying the booking limits for premium fare classes based on the capacities of each cabin. Figure 3-13 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.

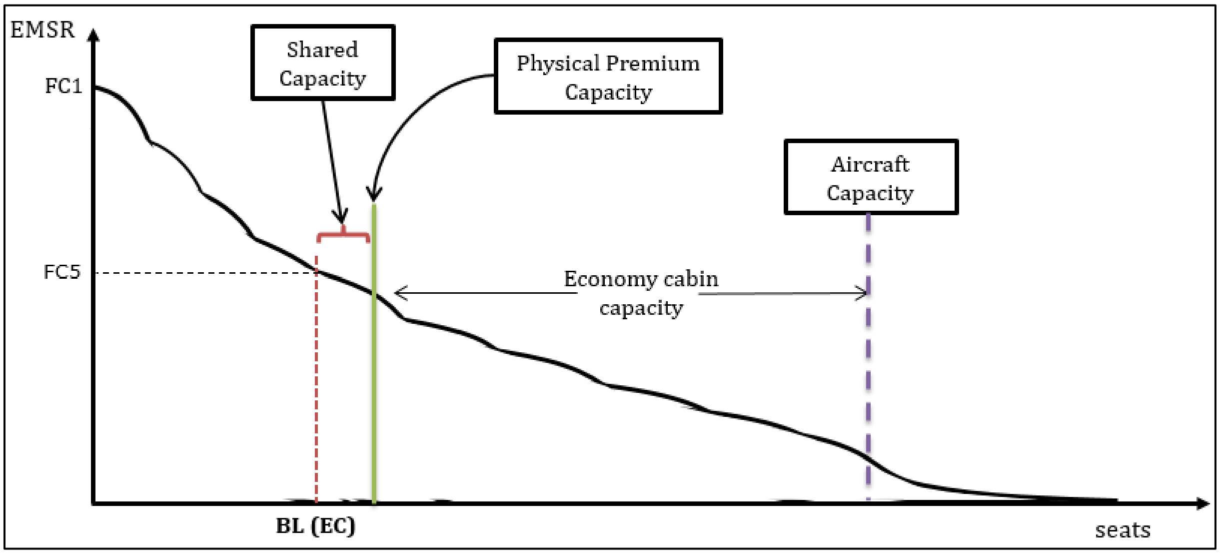

Figure 3-13: Full EMSR shared capacity

Table 3-13 provides an example of the differences between Distinct EMSRb and Full EMSR on a hypothetical leg operated by an aircraft with 20 seats in premium cabin and 150 seats in 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, it is determined that 16 seats have to be 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 initially 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 3-13: Distinct vs Full EMSR with shared capacity

It should be noted that Full EMSR does not necessarily lead to sharing premium capacity; as a matter of fact, the methodology ignores that passengers who purchase premium cabin fare classes cannot be accommodated in economy cabin. In such case, joint protection for premium fare classes has to be limited to the premium capacity and booking limits for economy fare classes must be increased. An example is presented in Table 3-14: 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 premium cabin. Hence, the booking limits for economy fare classes have to be adjusted in order to reflect the economy cabin capacity. 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.

3.5.2 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 and are used when determining if a booking request is accepted or not. In order to enable the premium cabin capacity sharing capability using a Bid Price Control mechanism (such as ProBP), the following schemes are proposed:

		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 3-14: Distinct vs Full EMSR without shared capacity

- **Premium Cabin Bid Price Capacity Sharing**

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 following conditions are met:

- Economy cabin is full in at least one of the legs in L_j.
- There are premium cabin seats available on all the legs in L_j.
- The fare $F_{i,j}$ of economy class i on ODF j is higher than sum of the premium cabin bid prices PBP_k of every leg k in L_j.

This mechanism applied for a single leg itinerary is illustrated in Figure 3-14. It should be clarified that when the conditions stated above are not met (e.g. there is economy cabin capacity available in all the legs in L_j), distinct cabin bid price control mechanisms described previously are still used.

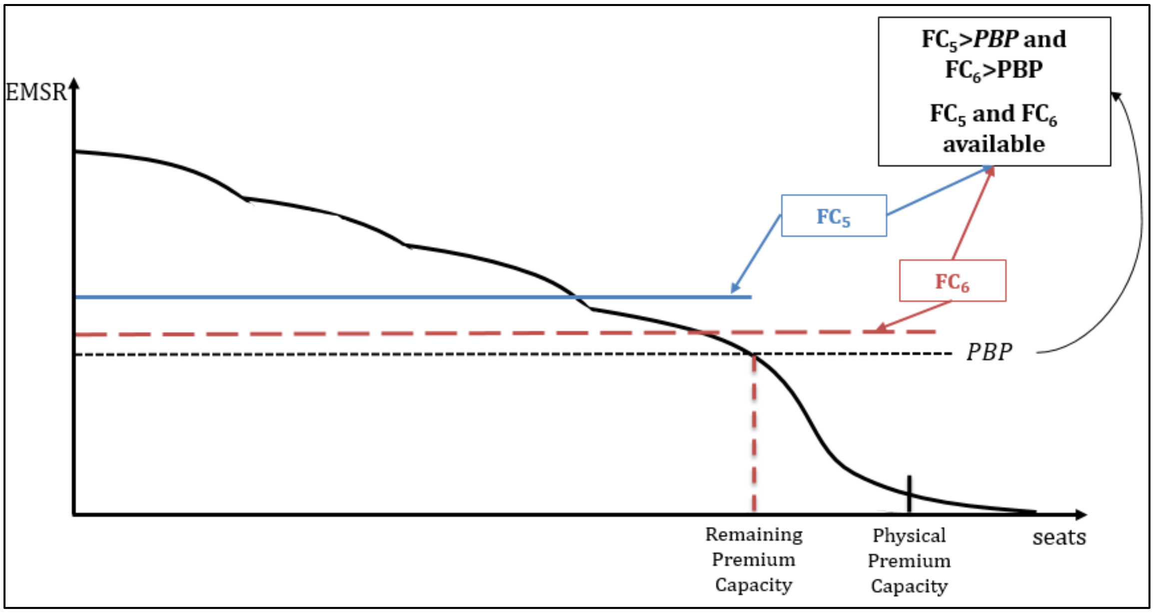

Figure 3-14: Premium cabin bid price control for a single leg itinerary

- **Joint Cabin Bid Price Capacity Sharing**

In this mechanism, a joint bid price JBP_k is calculated for each leg k considering the full remaining seating capacity; this means that its calculation ignores the division between premium and economy cabin. As with the previous bid price sharing mechanism this mechanism applies only when:

- Economy cabin is full in at least one of the legs in L_j.
- There are premium cabin seats available on all the legs in L_j.
- The fare $F_{i,j}$ of economy class i on ODF j is higher than sum of the joint cabin bid prices JBP_k of every leg k in L_j.

- **Effective Bid Price Capacity Sharing**

This scheme combines some of the features of each of the bid price control sharing mechanisms presented hereinabove; however, instead of applying only after a set of conditions are met, this mechanism applies through all the booking period. It requires the introduction of the Effective Premium Bid Price for leg k, $EPBP_k$, and of $EEBP_k$, the Effective Economy Bid Price, defined as follows:

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

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

In this case, bookings for fare classes are accepted if the fare is higher than the sum of the effective bid prices corresponding to the cabin of such fare class; this 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.

3.6 Summary

This chapter introduced the tools and methodologies that are used in the following sections of this document. First, the Passenger Origin Destination Simulator (PODS) was described and the revenue management optimizers used in this thesis were discussed: EMSRb for leg-based RM and ProBP for OD-control RM. To contextualize the discussion of the following chapters, the main characteristics of Network V1 and the baseline scenarios were presented. Then, the analytical methodology that is used in Chapter 4 to evaluate dual cabin aircraft configurations was introduced. Finally, the premium cabin capacity sharing mechanisms that are evaluated in Chapter 5 were described.

4. Cabin Configuration Analysis

This chapter examines the impact of changing the aircraft's layout of passenger accommodation (LOPA) on an airline's revenue and evaluates the performance of the application of the Boeing-Swan Spill Model ("BSM") to dual cabin aircraft by comparing its results against the outcome of simulations of such changes in PODS. The chapter is divided in four sections: the first section provides detailed information about the aircraft types considered for modification and the configuration changes proposed for each aircraft type. The second section describes the dual cabin BSM and presents the results of its application to the aircraft types identified previously. The chapter then explores the results for the simulations in PODS and compares these results with the ones obtained using the dual cabin BSM. Finally, Section 4.4 focuses on understanding the differences between the dual cabin BSM and the results of the simulations by examining the main drivers for such differences: load factor estimation and average spill fares.

4.1 Aircraft Types and Configuration Descriptions

As described in Chapter 3, Airline 1 in Network V1 operates a fleet that consists of seven different aircraft types; the configuration of each of these aircraft is divided into premium and economy cabins. Table 4-1 presents these aircraft types and provides the main operational parameters associated with each of them, as well as the share of AL1's revenue captured by each of these aircraft (connecting fares prorated based on mileage).

Corresponding to more than three fourths of AL1's total revenue, the three aircraft types with the higher capacities (i.e. 150-20, 200-24, 250-30) are selected in order to study the effect on revenue of certain configuration changes. Indeed, the selected aircraft types represent a good mix because of the characteristics of the legs operated by each of them; one of these characteristics is the proportion of domestic and international flights operated by each specific fleet type. As presented in Table 4-2, the vast majority of the legs operated by fleet type "A" (150-20) are between domestic city pairs. Fleet type "B" (200-24) operates half of its flights between domestic cities and the remaining half involves an international city either as the origin or the destination. Finally, all the legs operated by Airline 1 in fleet type "C" (250-30) depart from or arrive to an international city.

Aircraft Type (Economy Capacity-Premium Capacity)	Legs	Avg. Stage Length (mi)	Share of AL1's total revenue
50-9	4	549	0.7%
70-9	10	494	2.0%
100-12	32	1,049	9.8%
120-16	29	1,134	12.2%
150-20	30	1,271	18.0%
200-24	12	2,709	25.0%
250-30	9	3,668	32.3%

Table 4-1: Airline 1 Fleet

Aircraft Type	Domestic Legs	International Legs
A: 20-150	28 (93.3%)	2 (6.7%)
B: 24-200	6 (50%)	6 (50%)
C: 30-250	0 (0%)	9 (100%)

Table 4-2: Domestic and International Legs by Aircraft Type

It should be clarified that flying a domestic leg does not mean that all the passengers carried on such legs are travelling in a domestic OD pair: instead, some of them could be flying on those legs as part as an itinerary that involves connecting with an international leg. By contrast, all the passengers flying on an international leg are effectively travelling in an international market. This distinction could have an impact in the results considering that, on top of the disutility of traveling in economy cabin instead of premium, there is an additional disutility associated with traveling in economy cabin in international markets. Therefore, in addition to the passengers flying on international legs, those passengers flying on domestic legs that are connecting with an international leg are also subject to this disutility.

Another characteristic that differentiates the selected aircraft types is the proportion of local passengers carried on the legs operated by those aircraft. Table 4-3 indicates that the proportion of local passengers is higher on the legs operated by aircraft type 'A' but does not exceed 30% in any of the scenarios. As a matter of fact, the proportion of total passengers that fly an itinerary that involves connections is as high as 82% in some scenarios. These figures highlight the relevance of connecting passengers (i.e. non-local) in the network and the importance of considering network effects when analyzing any potential change to a portion of such network. Furthermore, this table also shows that the proportion of local passengers carried on those legs is typically higher when OD-control RM systems are in place.

	Aircraft Type	Local Passengers (Leg-Control)	Local Passengers (OD-Control)
Low	A: 150-20	27.7%	27.9%
	B: 200-24	24.3%	24.7%
	C: 250-30	17.7%	17.6%
Medium	A: 150-20	26.9%	27.9%
	B: 200-24	23.9%	25.2%
	C: 250-30	17.9%	18.0%
High	A: 150-20	25.9%	29.0%
	B: 200-24	22.9%	25.5%
	C: 250-30	18.0%	18.7%

Table 4-3: Proportion of local passengers by aircraft type and demand level

An additional feature of the legs operated by each aircraft type is related with the average leg load factor "ALLF" (Table 4-4 and Table 4-5 for leg-control and OD-control, respectively): Aircraft type A combines a considerably low ALLF (30%-43%, depending on the demand level) in its premium cabin with the highest economy ALLF among the fleet types (79%-91%). On the other hand, the flights operated by aircraft type C (100% international) exhibit the highest ALLF's in premium cabin among the evaluated aircraft types (78%-83%) combined with considerably high ALLF's in economy cabin (78%-86%). Finally, Aircraft type B has a moderate premium ALLF's (47%-60%, resulting from combining low premium LF's on domestic legs and relatively high premium LF's on international legs), with reasonably high economy ALLF's, in the range between 75% and 90%, depending on the demand level.

Demand	Aircraft Type	Baseline ALLF	Baseline Premium Cabin ALLF	Baseline Economy Cabin ALLF
Low	A: 150-20	73.39%	30.22%	79.15%
	B: 200-24	72.38%	47.27%	75.40%
	C: 250-30	78.28%	78.32%	78.27%
Medium	A: 150-20	79.68%	35.74%	85.54%
	B: 200-24	79.50%	53.18%	82.65%
	C: 250-30	81.98%	80.88%	82.11%
High	A: 150-20	85.71%	43.39%	91.35%
	B: 200-24	86.36%	59.66%	89.57%
	C: 250-30	85.84%	83.47%	86.12%

Table 4-4: Average Leg Load Factor (ALLF) by aircraft type and demand level – Leg Control

Demand	Aircraft Type	Baseline ALLF	Baseline Premium Cabin ALLF	Baseline Economy Cabin ALLF
Low	A: 150-20	72.65%	30.30%	78.29%
	B: 200-24	73.36%	48.27%	76.38%
	C: 250-30	78.67%	77.82%	78.78%
Medium	A: 150-20	77.08%	35.67%	82.60%
	B: 200-24	79.86%	53.85%	82.98%
	C: 250-30	83.96%	81.12%	84.30%
High	A: 150-20	81.50%	41.85%	86.79%
	B: 200-24	85.10%	59.39%	88.19%
	C: 250-30	88.43%	84.21%	88.93%

Table 4-5: Average Leg Load Factor (ALLF) by aircraft type and demand level – OD-Control

As discussed in Chapter 3, four alternate configurations are evaluated for each of these aircraft types assuming that adding (or removing) a row of premium cabin seats results in the removal (or addition) of two rows of economy cabin seats. The evaluated configurations are presented in Table 4-6. When premium cabin seats are added to the aircraft configuration, the total capacity of the aircraft is reduced; on the other hand, replacing premium cabin seats with economy cabin seats results in an increase of total capacity. Therefore, the capacity offered by Airline 1, measured in

ASMs, changes as a result of these configuration changes; Table 4-7 presents the proportional changes in the ASMs of Airline 1.

A/C Type	Configuration High Premium (HP)	Configuration More Premium (MP)	Configuration Baseline (BL)	Configuration More Economy (ME)	Configuration High Economy (HE)
A	126-28	138-24	**150-20**	162-16	174-12
B	168-36	184-30	**200-24**	216-18	232-12
C	218-42	234-36	**250-30**	266-24	282-18

Table 4-6: Aircraft Seat Configurations Simulated

A/C Type	Configuration High Premium (HP)	Configuration More Premium (MP)	Configuration Baseline (BL)	Configuration More Economy (ME)	Configuration High Economy (HE)
A	-1.92%	-0.96%	-	0.96%	1.92%
B	-2.08%	-1.04%	-	1.04%	2.08%
C	-2.05%	-1.02%	-	1.02%	2.05%

Table 4-7: Proportional Variation in ASM's in the network compared to the baseline

4.2 Analytical Method for Evaluating Configuration Changes

As presented in Chapter 3, the analytical method used to estimate the impact of a change of the capacity of the aircraft operating a leg on the revenue captured from the leg is the Boeing-Swan Spill Model ("BSM"). This model considers the effect on the passenger loads resulting from reducing or increasing the capacity offered on a leg. Then, it captures the effect over the airline's revenue by assigning an estimated fare to those spilled (when capacity is reduced) or additionally accommodated (when capacity is increased) passengers.

4.2.1 Adjustment of the BSM for Dual Cabin Aircraft

The single cabin aircraft BSM considers a flight that has been operated with an aircraft with capacity C_1, in which a historical load factor L_1 has been observed; using the formulation presented in section 3.4.2, a demand factor D_1 can be estimated for such flight leg by assuming a factor K. This demand factor D_1 allows for the calculation of the unconstrained demand μ for that leg. By modifying the capacity to be C_2, a new demand factor D_2 can be calculated; L_2 can be then determined from the model (or obtained from tables provided by Boeing), and the expected loads $L_2 \times C_2$ can be calculated.

For the case of dual cabin aircraft analyzed in this thesis, increasing the capacity of one of the cabins reduces the number of seats available in the other cabin. Therefore, the model is implemented simultaneously in each of the cabins: on the one hand it is used to estimate the additional passengers that will be carried in the cabin that has its size increased; on the other hand, it calculates the number of passengers spilled in the cabin with reduced capacity. This adaptation of the BSM model for dual cabin aircraft is hereinafter referred to as the dual cabin BSM.

Some of the assumptions made by this approach are discussed herein in order to provide insights for analyzing the results presented in the following sections. One of such assumptions is that the model considers that each cabin has its own unconstrained demand (say, μ_P for premium and μ_E for economy), independent from the other cabin. This implies that a passenger would not consider flying in a cabin different from the one originally requested even if an option available in the other cabin provides higher utility (e.g. an "economy cabin passenger" would not consider at all buying a ticket in premium cabin). Closely related with the first assumption, the model does not consider that passengers spilled from a cabin represent additional demand for travel on the other cabin; similarly, it ignores that providing additional capacity in a cabin could result in passengers deciding to book their ticket in such cabin instead of the cabin where the booking request would have been received should the capacities have not been modified.

Furthermore, there are other assumptions that are shared by the single cabin BSM and the dual cabin BSM that are worth some discussion. The first of these is that while the BSM model describes the unconstrained demand as coming from a single source (passengers flying specifically between the origin and the destination of the flight leg), the reality is that the passengers on a flight are typically travelling between a diverse group of OD markets, using that specific leg just as a portion of their trip. In addition, the BSM ignores any interaction between flights operated by an airline on the same route: for example, some passengers spilled from a flight that has its capacity reduced could possibly decide to book on another flight operated by the airline that covers the same route at another time of the day; this limitation extends to the lack of consideration by the BSM for any form of competition.

The assumptions discussed so far are mostly related with the estimation of the carried loads on the flight leg under analysis; however, this information has to be combined with the estimation of the passenger value in order to measure the impact of the on an airline's revenue. As discussed in Chapter 3 and following the BSM, the proposed average fare of a spilled passenger considered in this thesis is a weighted average where 80% corresponds to the lowest fare in the market and the remaining 20% to the average fare in that same market. However, there is a slight modification that needs to be taken into consideration in order to adjust the model to the fare structure of an airline flying dual cabin aircraft: for the premium cabin, the weighted average is between the average fare for the premium fare classes and the lowest fare among that same group (i.e. FC4). Likewise, for the economy cabin the considered fare is a combination of the average and lowest fare specifically applicable for economy cabin.

Two other considerations related with fares are discussed herein: first, it should be noted that the same fares are used when the number of seats in the cabin increases or decreases. However, when RM effects are taken into consideration, the magnitude of the revenue loss of a spilled passenger should be higher or equal to the gain of an additional passenger because it is assumed that higher value passengers are on the flight instead of passengers of lower value. Finally, it should also be noted that, since there are connecting passengers flying on the legs being analyzed, the average and lowest fare are not the fares of the specific OD market covered by the leg, but the combination of the prorated fares for all the passengers flying on such leg.

4.2.2 Dual Cabin Boeing Spill Model: Results

- **Aircraft Type 'A' (baseline Economy Cabin 150 seats – Premium Cabin 20 seats)**

Figure 4-1 presents the proportional variation in AL1's total revenue (premium + economy) estimated using the dual cabin BSM, for different LOPA configurations applicable to aircraft type A. The pattern is clear: AL1 benefits from operating aircraft type A in a configuration with more economy cabin seats. This result is expected when the cabin ALLF's for each cabin, referenced in section 4.1.1, are considered together: while high economy cabin ALLF's suggest that the baseline economy cabin capacity could be the factor limiting the amount of bookings accepted by AL1 for economy fare classes and therefore additional economy cabin capacity is beneficial, the low premium cabin ALLF's indicate that increasing the premium cabin capacity should not have a considerable impact on the loads carried. Furthermore, the revenue gains resulting from providing additional economy cabin capacity increase as the demand increases.

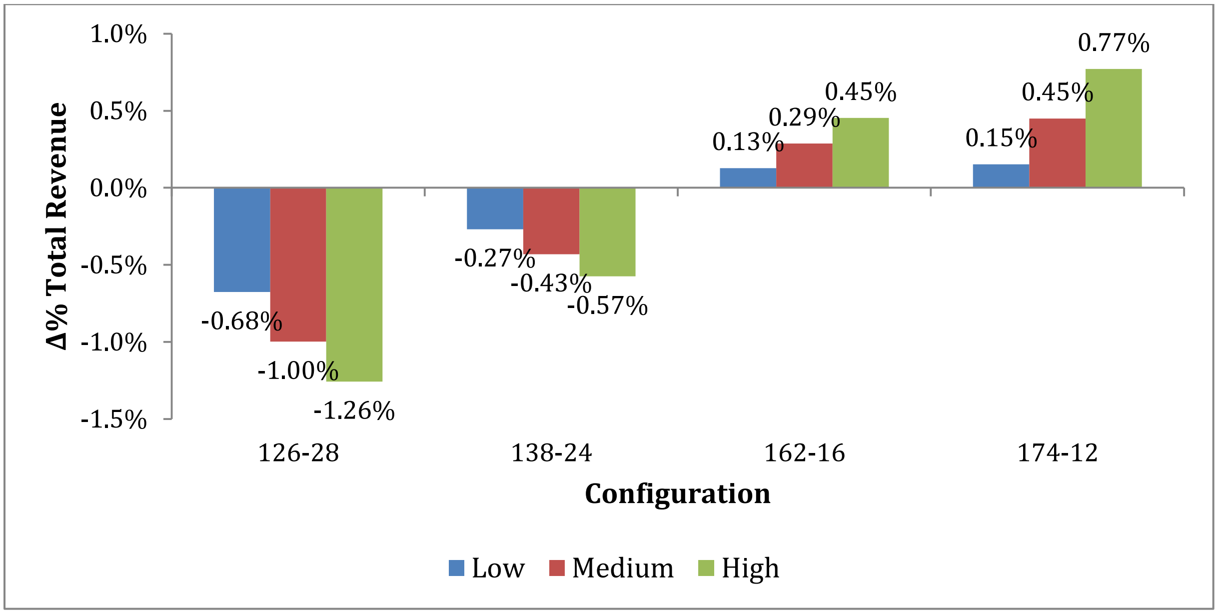

Figure 4-1: BSM Model - Total Revenue Proportional Variation – Aircraft Type A - Leg-RM

It should be noted that the figure shows the variation of the total revenue of AL1; however, as discussed in the preceding section, aircraft type A's share of AL1's total revenue is just 18.3%. Since the BSM considers only changes in revenue on the legs effectively operated by the fleet, all the changes in revenue are concentrated on those 30 legs; the revenue variation specific to such legs is illustrated in Figure 4-2. A similar pattern is observed (revenue gains when more economy cabin seats are made available), but the magnitude of the proportional changes increases substantially.

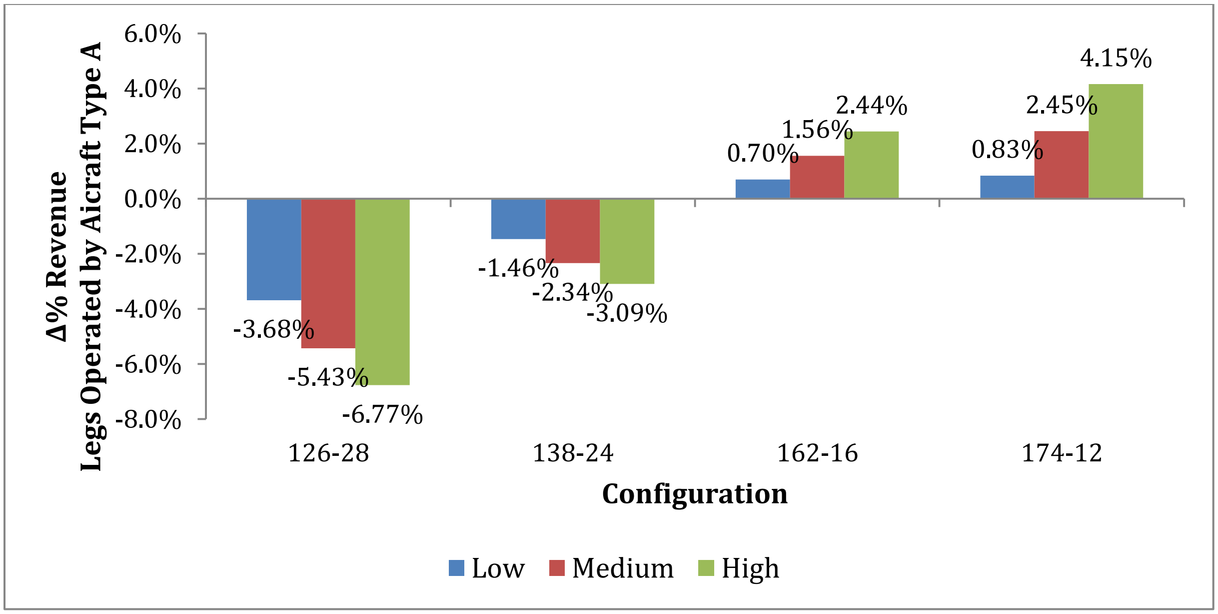

Figure 4-2: BSM Model – Proportional Variation Revenue on legs operated by A/C Type A - Leg-RM

The pattern of the revenue changes is similar for the OD-control RM method scenarios (Figure 4-3); however, the proportional changes are typically smaller in magnitude than in the leg RM scenario. Since the BSM does not consider any network effects at all, the differences between leg-control and OD-control are entirely based on characteristics of the baseline scenario used for each case. More specifically, since economy cabin ALLF's are lower in the OD-control baseline than in the leg-control case, there is less pressure for additional economy cabin seats, resulting in slightly lower revenue gains when additional capacity is made available at the medium and high scenarios (or losses, when the economy cabin capacity is reduced).

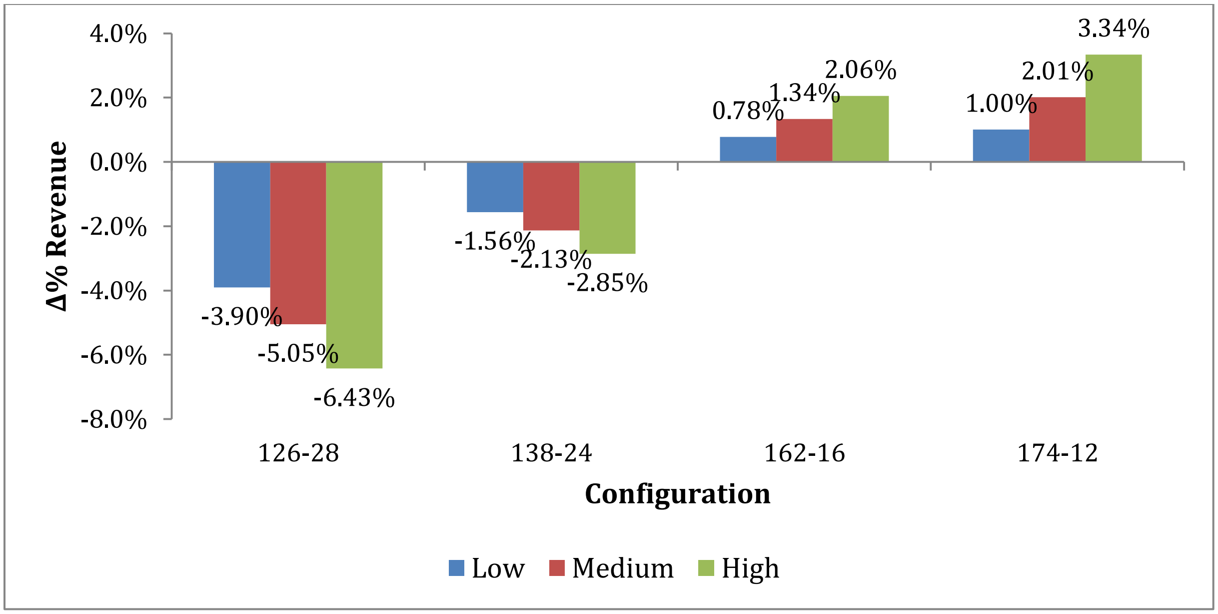

Figure 4-3: BSM Model - Proportional Variation Revenue on legs operated by A/C Type A – OD-Control

- **Aircraft Type 'B' (baseline Economy Cabin 200 seats – Premium Cabin 24 seats)**

The revenue changes estimated by using the dual cabin BSM for evaluating configuration changes to aircraft type B suggest that the baseline configuration (200-24) is the most adequate for the legs operated by such aircraft. As illustrated in Figure 4-4, the aggregated revenue of the twelve legs operated with this aircraft type decreases in almost all the alternative arrangements proposed (with the sole exception of the 216-18 configuration in the high demand scenario, where a revenue increase is estimated). Besides, configurations that increase premium cabin capacity hurt AL1's revenues more than the LOPA's that add economy cabin seats, mostly at the high demand scenarios.

However, as discussed earlier in this chapter, half of the routes operated by aircraft type B in AL1's network are international and half are domestic; Table 4-8 illustrates how ALLF's differ between these groups. While economy cabin ALLF's are considerably high for both cases, the premium cabin ALLF's are roughly 30 percentage points higher on the international legs. Therefore, different impacts are expected for each of these categories, as confirmed by Figure 4-5: while the 200-24 LOPA is the best suited for the international legs operated by aircraft type B, the model estimates that domestic legs would benefit from replacing premium cabin seats with economy cabin seats (similar to what has been observed for aircraft type A).

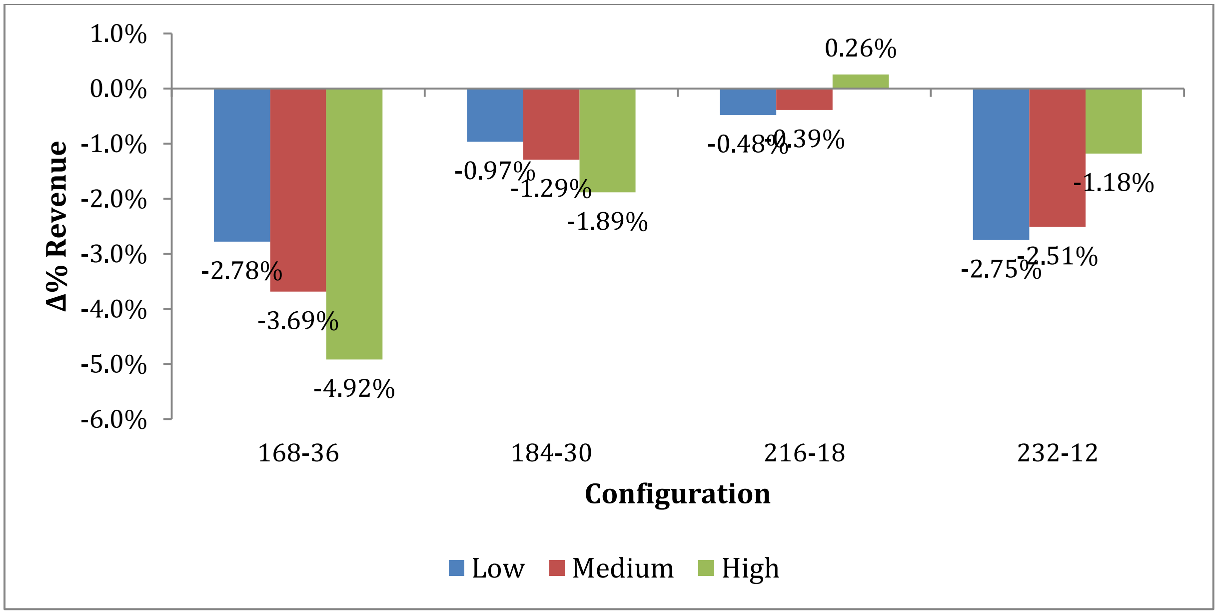

Figure 4-4: BSM Model – Proportional Variation Revenue in legs operated by A/C Type B – Leg-RM

Category	Demand	Premium Cabin ALLF	Economy Cabin ALLF
International	Low	61.83%	77.80%
	Medium	68.13%	83.78%
	High	73.32%	88.17%
Domestic	Low	32.72%	73.00%
	Medium	38.24%	81.52%
	High	46.00%	90.97%

Table 4-8: Average Leg Load Factors by leg category - Aircraft Type B – Leg-RM

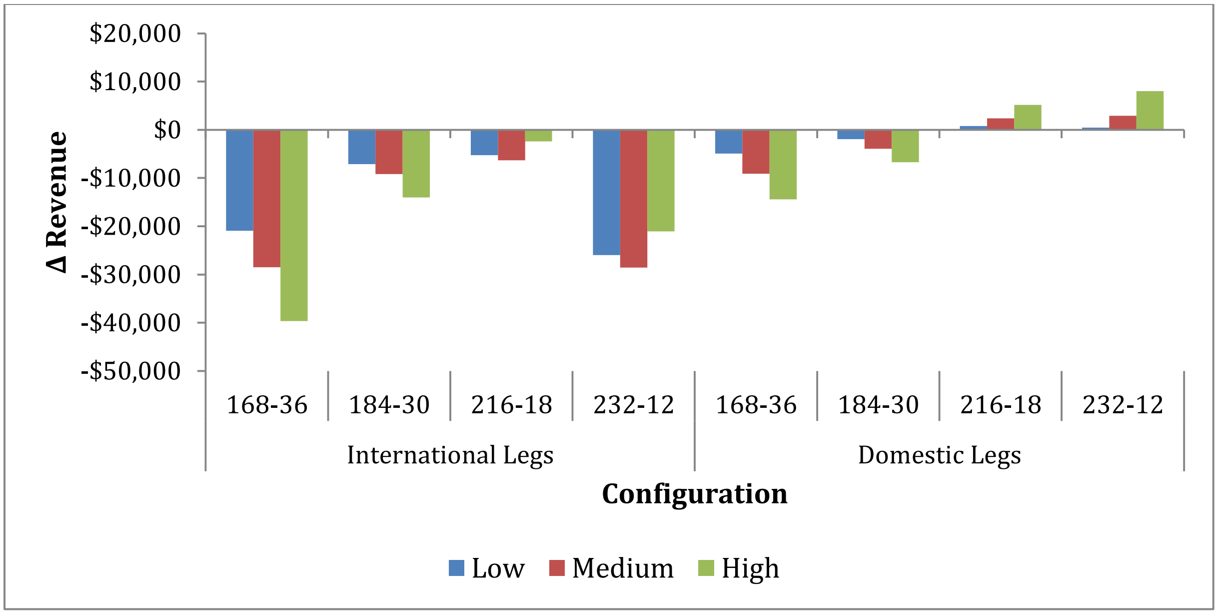

Figure 4-5: BSM Model - Absolute Revenue Variations by Leg Category - Aircraft Type B - Leg-RM

The proportional revenue variations are also presented for the OD-control baseline; again, the pattern is very similar to the leg RM baseline since the dual cabin BSM does not incorporate any network consideration in the estimation of the loads. Therefore, the differences between the magnitudes of the changes presented in Figure 4-6 (when compared to the leg-control baseline) are explained mainly by the differences in the economy cabin ALLF's on each of such baselines.

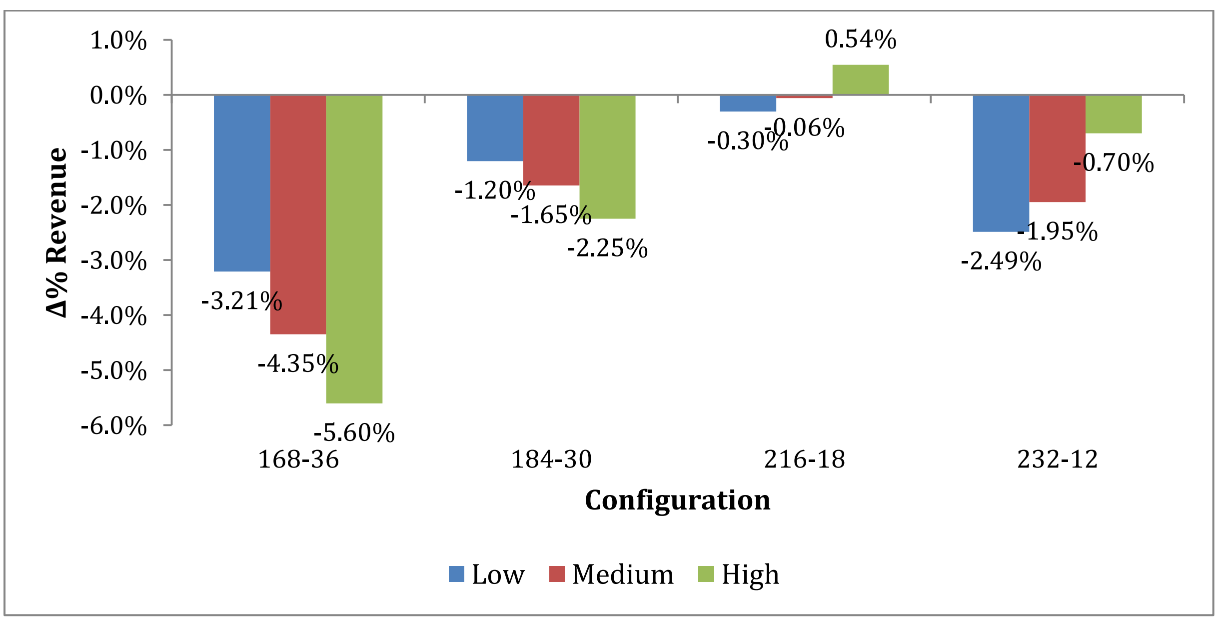

Figure 4-6: BSM Model – Proportional Variation Revenue on legs operated by A/C Type B - OD-control

- **Aircraft Type 'C' (baseline Economy Cabin 250 seats – Premium Cabin 30 seats)**

The proportional total revenue changes with respect to the baseline configuration are also estimated using the dual cabin BSM for aircraft type C (Figure 4-7). In this case, highest total revenues are observed in the configurations in which the number of premium cabin seats is increased (i.e. 218-42 and 234-36). However, the revenue gains in such configurations decrease as demand increases, since ALLF's in the baseline configuration are also high for economy cabin and the reduced economy cabin capacity leads to a large number of economy cabin passengers spilled when demand is high. When the economy cabin capacity increases, the total revenue decreases because of the premium cabin spill; nevertheless, as demand increases and the number of additional economy cabin passengers increases consequently, the total revenue losses are reduced.

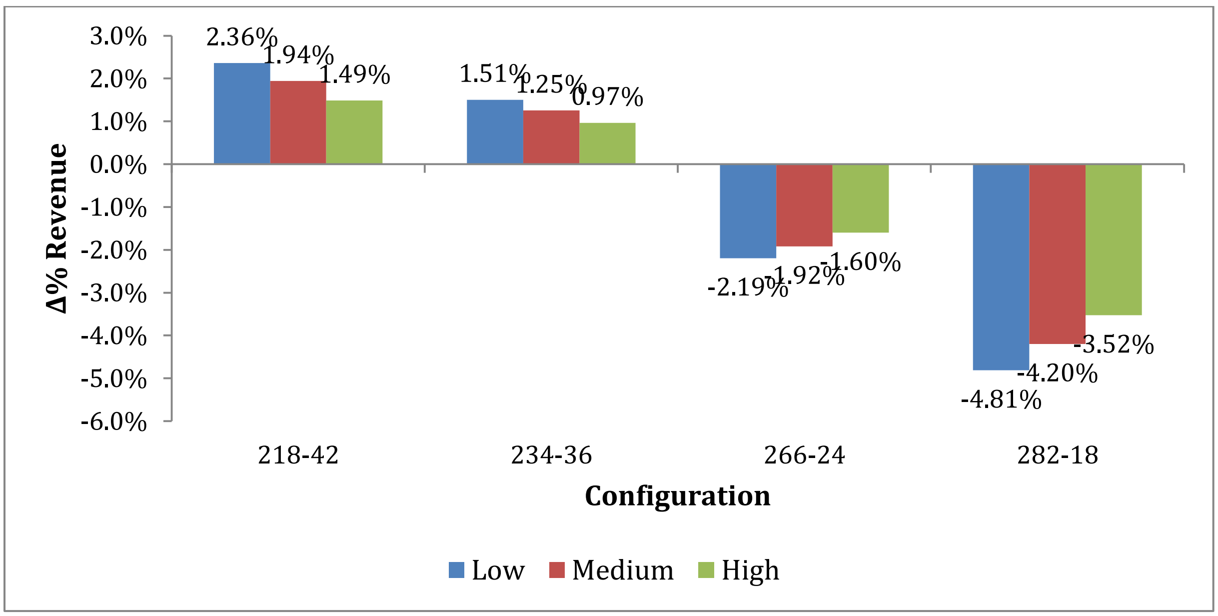

Figure 4-7: BSM Model – Proportional Variation Revenue in legs operated by A/C Type C - Leg-RM

In the OD-control baseline the pattern is similar (although the magnitude of the proportional changes is lower than in the leg-control baseline) for the low and medium demand scenarios. However, the cabin configurations with highest premium cabin capacities reduce their revenue at the high demand scenario (Figure 4-8). As found for the other aircraft types, the high economy cabin ALLF observed in the baseline configuration at the high demand (88.93%) leads to a large number of economy cabin passengers spilled that are not compensated by the number of additional passengers carried in premium cabin (again, it should be emphasized that to add six premium cabin seats, sixteen economy cabin seats must be removed).

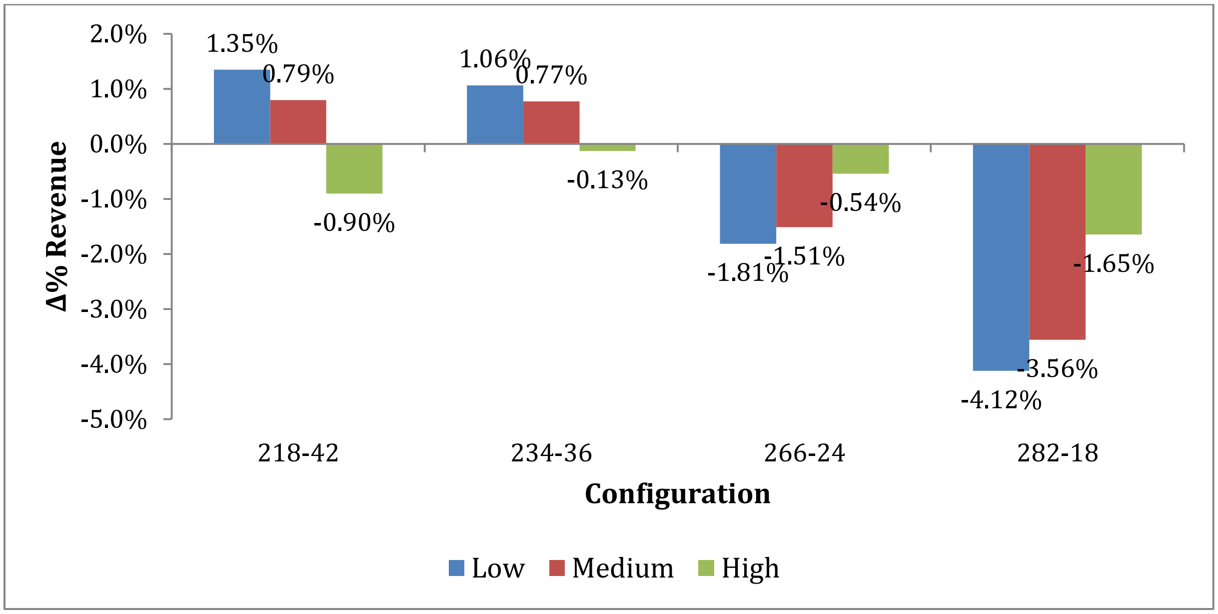

Figure 4-8: BSM Model – Proportional Variation Revenue in legs operated by A/C Type C - OD-control

4.3 Simulations

As discussed in Section 4.2, the adjusted version of the BSM for dual cabin aircraft relies on many assumptions. For example, it assumes independent demands for each cabin, independent effects of the leg with respect to the rest of the network and lack of competition. In order to address all these limitations of the BSM, each of the configuration changes evaluated herein are simulated in PODS at three different demand levels. The results of these simulations are analyzed in this section.

- **Aircraft Type 'A' (baseline Economy Cabin 150 seats – Premium Cabin 20 seats)**

Figure 4-9 presents the total revenue (premium + economy) changes for AL1 resulting from modifying the LOPA of aircraft type "A" when leg RM is used. Overall, the proportional changes follow the same pattern identified with the dual cabin BSM in section 4.2 for this same aircraft type: total revenue losses when premium cabin capacity increases and gains when additional economy cabin capacity is made available (Figure 4-1). At first sight, the results obtained from the simulations are similar to the ones estimated using the dual cabin BSM; nevertheless, the simulations in PODS allow us to examine the effect of the configuration changes not only on the legs operated by the aircraft type under analysis but also on the rest of the network.

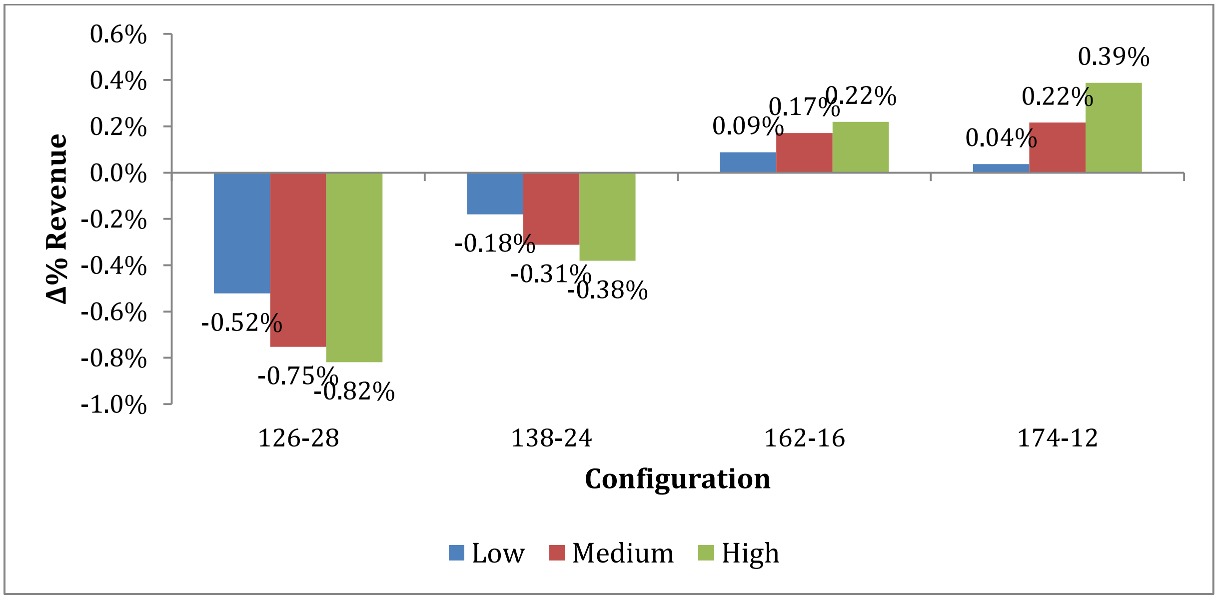

Figure 4-9: PODS - Total Revenue Proportional Variation – Aircraft Type A - Leg-control

Figure 4-10 provides a breakout of the proportional revenue changes for two sets of legs: those operated by aircraft type A and those operated by the rest of AL1's fleet (keeping the configuration of the rest of the fleet unchanged). As discussed above, the magnitude of the proportional variations is larger when the affected legs are examined apart from the rest of the network; however, such variations are not as large as predicted by the dual cabin BSM. For example, while the spill model estimated proportional losses of -6.77% with the 126-28 configuration at the high demand level (Figure 4-2), the outcome of PODS indicates a reduction in the revenue of such legs of -4.89%; in this specific case, both gains and losses are exaggerated by the dual cabin BSM and such differences are exacerbated at high demand levels, as summarized in Table 4-9.

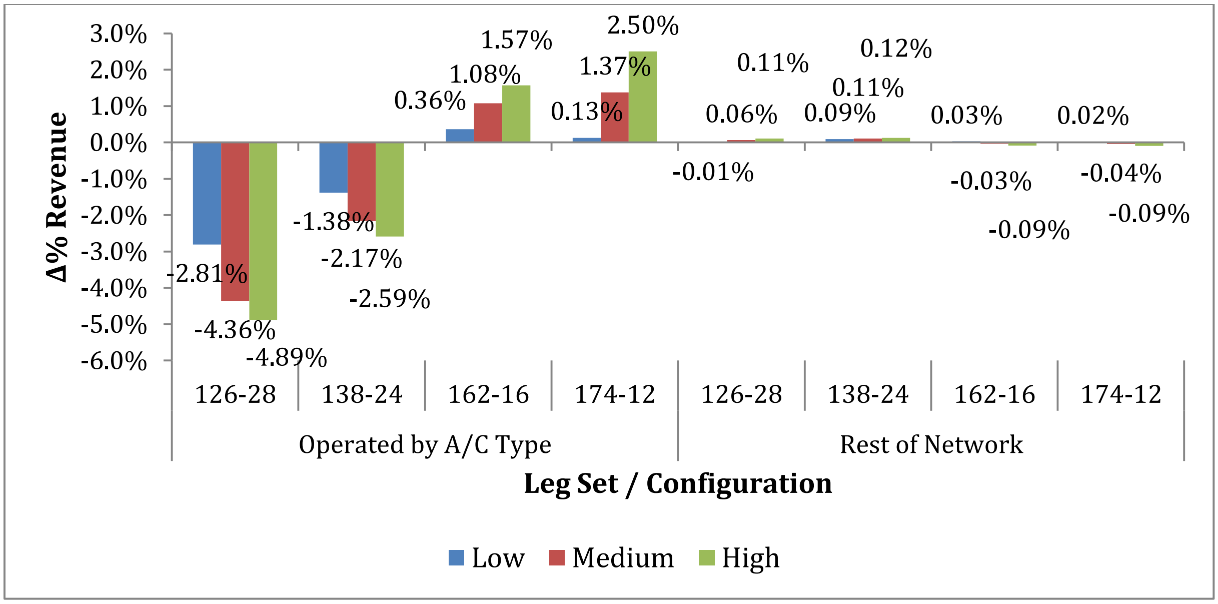

Figure 4-10: PODS - Proportional Revenue Variation by set of legs – Aircraft Type A – Leg Control

		Configuration			
Demand	**Mechanism**	**126-28**	**138-24**	**162-16**	**174-12**
Low	BSM	-3.68%	-1.46%	0.70%	0.83%
	PODS	-2.81%	-1.38%	0.36%	0.13%
Medium	BSM	-5.43%	-2.34%	1.56%	2.45%
	PODS	-4.36%	-2.17%	1.08%	1.37%
High	BSM	-6.77%	-3.09%	2.44%	4.15%
	PODS	-4.89%	-2.59%	1.57%	2.50%

Table 4-9: Proportional revenue variations on legs operated by aircraft A – Leg-based RM

In addition to the differences between the PODS results and the dual cabin BSM model for the legs operated by aircraft A, Figure 4-10 also illustrates how the rest of the network is affected by the changes made on a subset of the legs; although the proportional variations are low (less than 0.12%) it shows that the rest of the network benefits mostly from configurations of aircraft type A with the highest premium cabin capacities, as opposed to what has been found for the legs operated with the modified aircraft. Moreover, since more than 81% of AL1's revenue is captured from legs different from the ones operated by aircraft type A, the absolute revenue variations observed in the rest of the network represent a significant portion the total revenue variations and help to compensate partially for the variations on the legs (Figure 4-11).

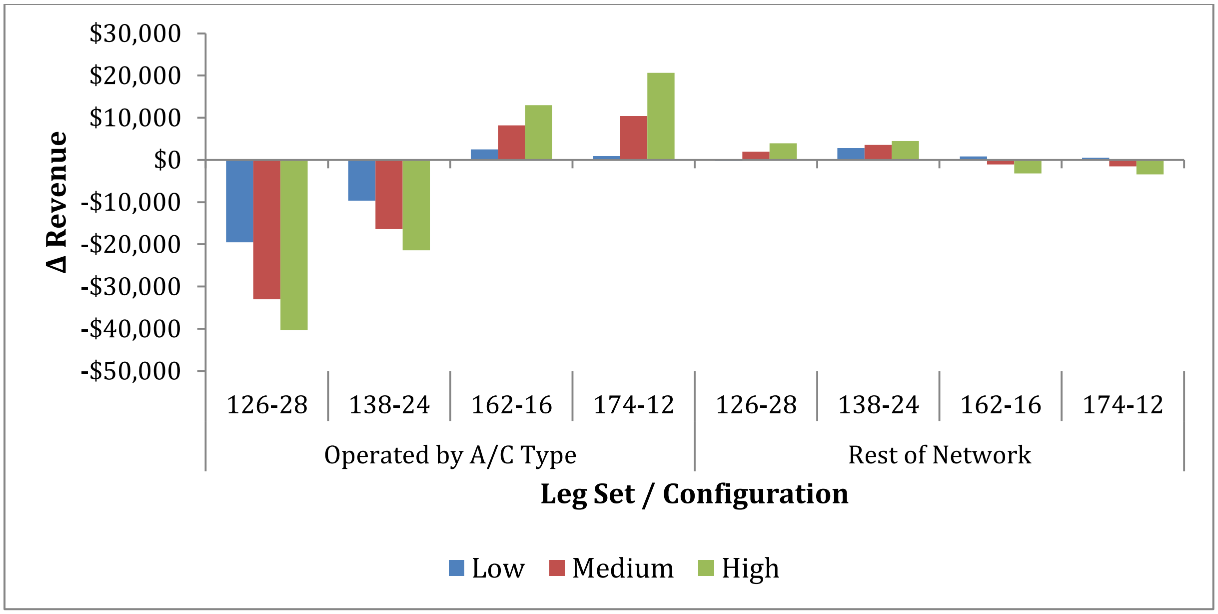

Figure 4-11: PODS - Absolute Revenue Variation by set of legs – Aircraft Type A – Leg Control

Finally, the proportional revenue variations are presented in Figure 4-12 for the OD-control baseline. Compared with the variations estimated using the dual cabin BSM, it is found again that the simulated variations are smaller in magnitude in most of the scenarios. In addition, the magnitudes of the proportional variations observed for the rest of the network are also larger with OD-control than with leg-RM (in most of the cases). Therefore, those other legs make a substantial contribution to AL1's total revenue when additional premium capacity is made available in aircraft type A (Figure 4-13). As in the case of leg-based RM, the BSM is able to determine if the configuration change would result in revenue gains or losses for AL1. However, it also tends to exaggerate the gains and the losses resulting from a configuration change, with the exception of some scenarios in which the magnitude of the gains/losses is smaller in the BSM (light gray in Table 4-10).

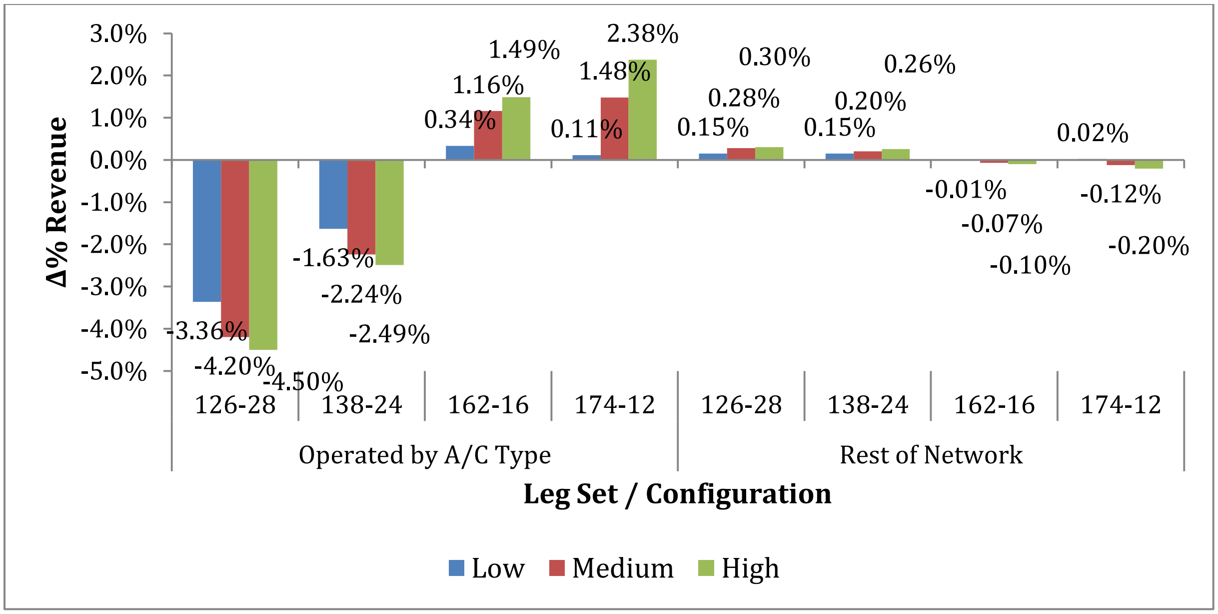

Figure 4-12: PODS - Proportional Revenue Variation by set of legs – Aircraft Type A – OD-Control

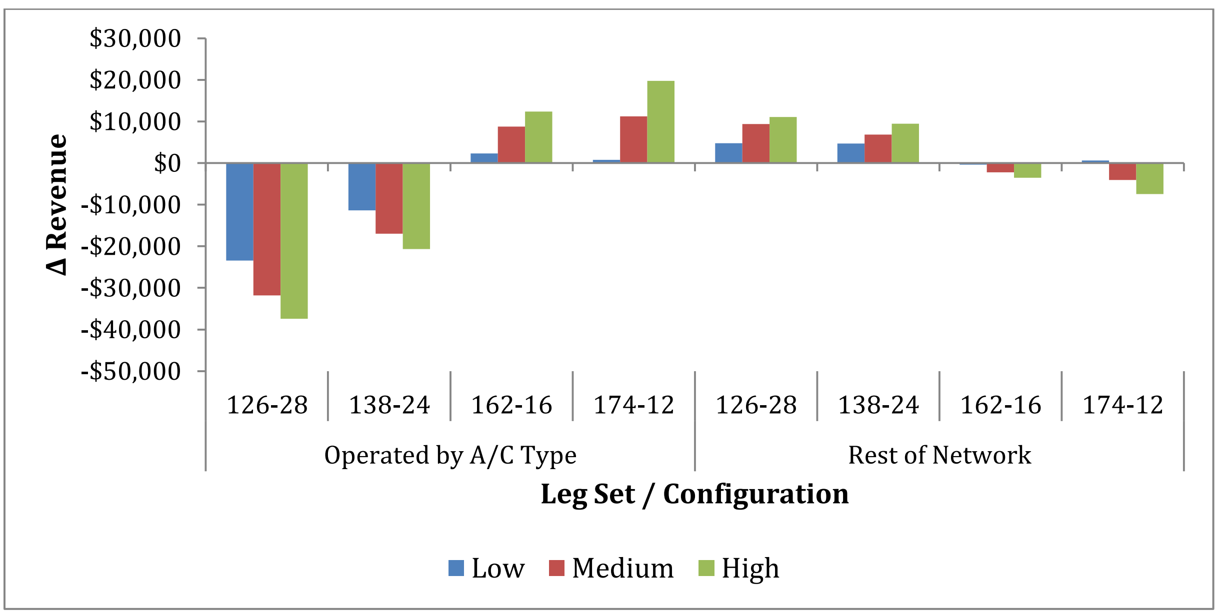

Figure 4-13: PODS - Absolute Revenue Variation by set of legs – Aircraft Type A – OD Control

		OD Control			
Demand	**Mechanism**	**126-28**	**138-24**	**162-16**	**174-12**
Low	BSM	-3.90%	-1.56%	0.78%	1.00%
	PODS	-3.32%	-1.59%	0.38%	0.15%
Medium	BSM	-5.05%	-2.13%	1.34%	2.01%
	PODS	-4.34%	-2.39%	1.01%	1.33%
High	BSM	-6.43%	-2.85%	2.06%	3.34%
	PODS	-3.90%	-1.88%	2.13%	3.02%

Table 4-10: Proportional revenue variations on legs operated by aircraft A – Leg-based RM

- **Aircraft Type 'B' (baseline Economy Cabin 200 seats – Premium Cabin 24 seats)**

The results of the simulations in PODS for the modification of aircraft type B confirm what was found with the dual cabin BSM: since the revenue decreases when any other configuration is used (at all demand levels), the 200-24 configuration seems to be the best configuration for the legs operated by such aircraft type (Figure 4-14). Nevertheless, it is also found that the BSM model tends to exaggerate the proportional revenue variations.

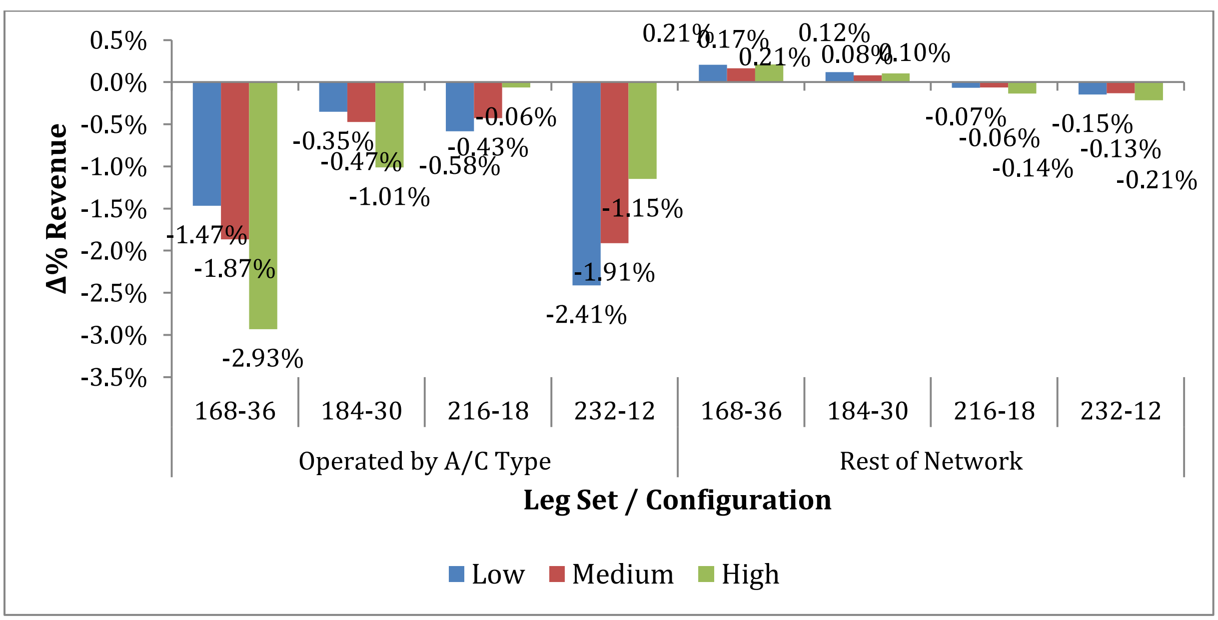

Figure 4-14: PODS - Proportional Revenue Variation by set of legs – Aircraft Type B – Leg Control

These simulations also confirm the effect over the rest of the network of changing the configuration of the aircraft operating a subset of the legs. Not only are the proportional variations higher than the ones presented in the case of aircraft type A (higher proportion of connecting passengers on legs operated by aircraft type B), but the impact over the total revenue variations becomes higher as well; as a matter of fact, in some cases the revenue variations in the rest of the network are higher than on the legs where the capacity was modified (illustrated in Figure 4-15). This graph also confirms that the rest of the network benefits from aircraft type B being configured with a LOPA with a large number of premium cabin seats.

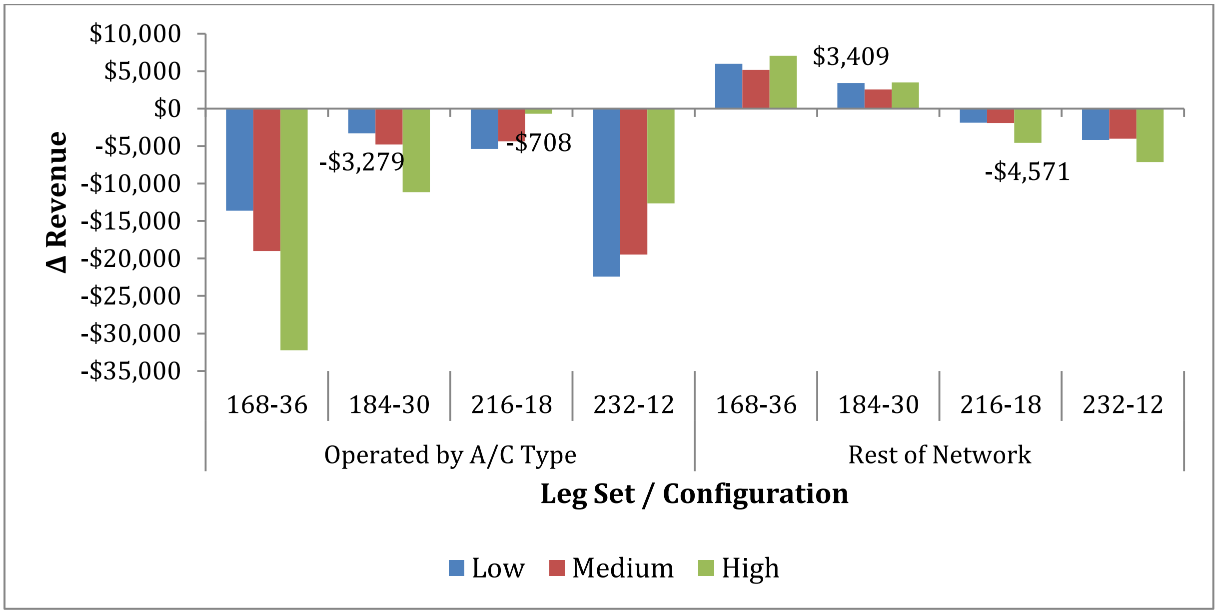

Figure 4-15: PODS - Absolute Revenue Variation by set of legs – Aircraft Type B – Leg Control

Finally, the proportional variations for the OD-control baseline are presented in Figure 4-16. These variations are consistent with the findings in the scenarios analyzed previously: proportional revenue variations that are lower in the simulations than in the dual cabin BSM for the legs operated by the aircraft modified.

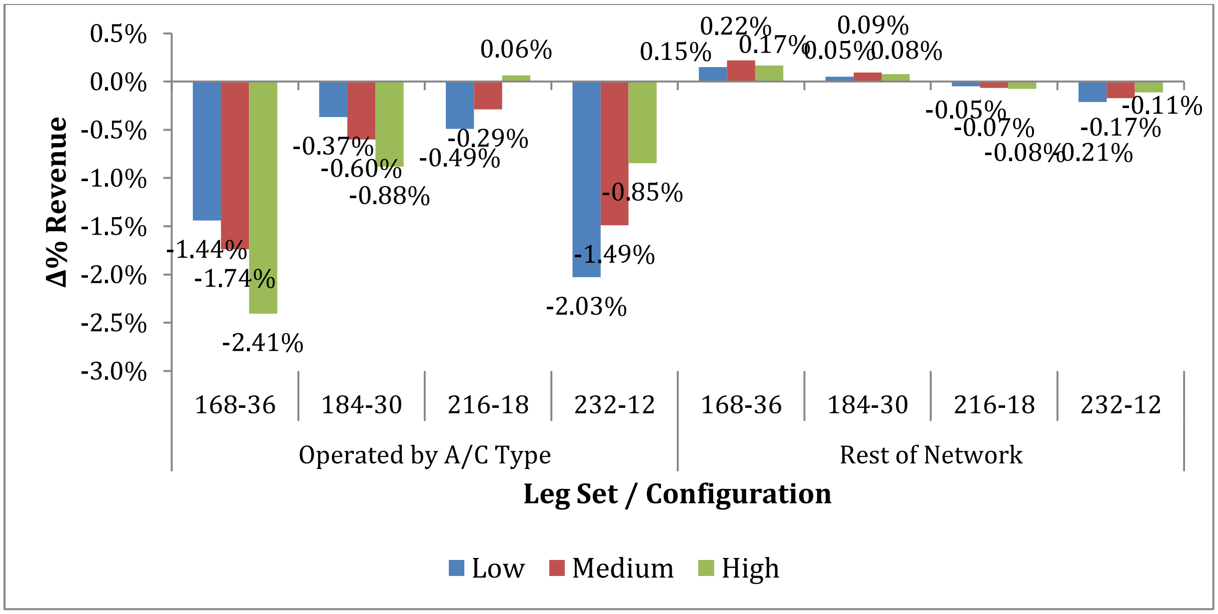

Figure 4-16: PODS - Proportional Revenue Variation by set of legs – Aircraft Type B – OD-Control

Table 4-11 and Table 4-12 present the proportional revenue variations with respect to the baseline configuration for BSM and PODS, compared to the baseline configuration of Aircraft B (200-24). As in the case of aircraft type A, in most of the cases the BSM overestimates the magnitudes of the gains or losses; however, there are some cases where the BSM estimation is smaller than in the simulations.

		Configuration			
Demand	**Mechanism**	**168-36**	**184-30**	**216-18**	**232-12**
Low	BSM	-2.78%	-0.97%	-0.48%	-2.75%
	PODS	-1.47%	-0.35%	-0.58%	-2.41%
Medium	BSM	-3.69%	-1.29%	-0.39%	-2.51%
	PODS	-1.87%	-0.47%	-0.43%	-1.91%
High	BSM	-4.92%	-1.89%	0.26%	-1.18%
	PODS	-2.93%	-1.01%	-0.06%	-1.15%

Table 4-11: Proportional revenue variations on legs operated by aircraft B – Leg-based RM

		Configuration			
Demand	**Mechanism**	**168-36**	**184-30**	**216-18**	**232-12**
Low	BSM	-3.21%	-1.20%	-0.30%	-2.49%
	PODS	-1.44%	-0.37%	-0.49%	-2.03%
Medium	BSM	-4.35%	-1.65%	-0.06%	-1.95%
	PODS	-1.74%	-0.60%	-0.29%	-1.49%
High	BSM	-5.60%	-2.25%	0.54%	-0.70%
	PODS	-2.41%	-0.88%	0.06%	-0.85%

Table 4-12: Proportional revenue variations on legs operated by aircraft B – OD-Control

- **Aircraft Type 'C' (baseline Economy Cabin 250 seats – Premium Cabin 30 seats)**

The results of the simulations for the scenarios in which aircraft type C is modified also confirm what was found with the dual cabin BSM, at least directionally: the legs operated by this aircraft would increase their revenue if additional premium cabin seats were made available. Figure 4-17 presents the proportional variations for the leg-control baseline and Figure 4-18 does the same for the OD-control baseline. In addition, the rest of the network increases its revenue when more premium cabin seats are made available in the legs operated by aircraft C.

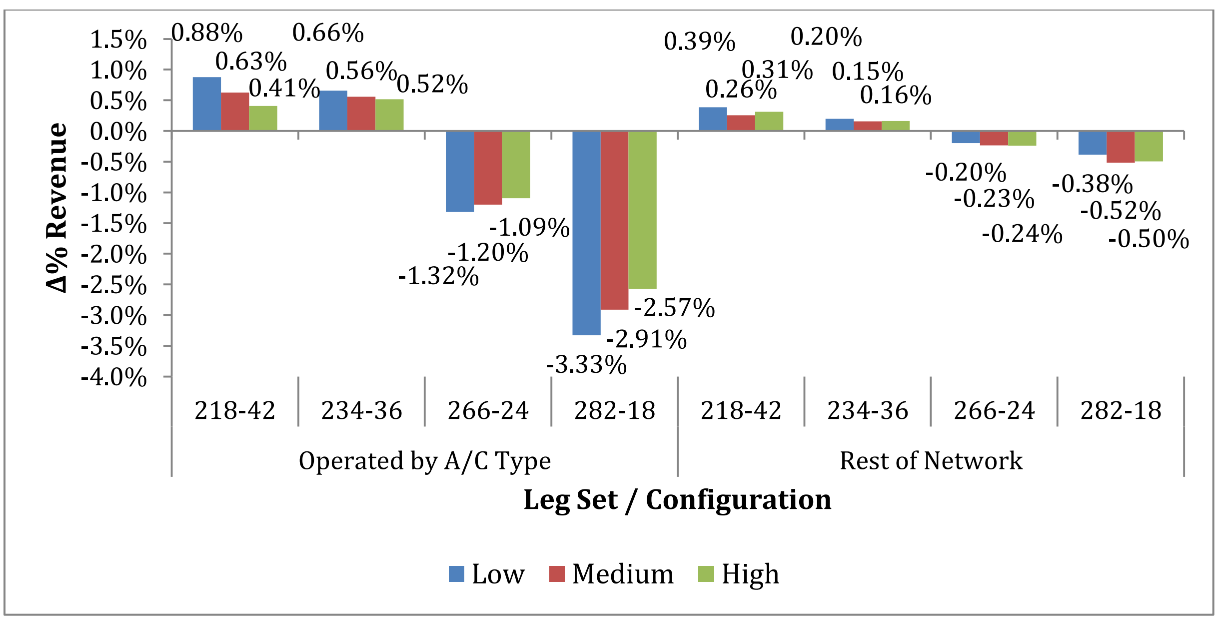

Figure 4-17: PODS - Proportional Revenue Variation by set of legs – Aircraft Type C – Leg Control

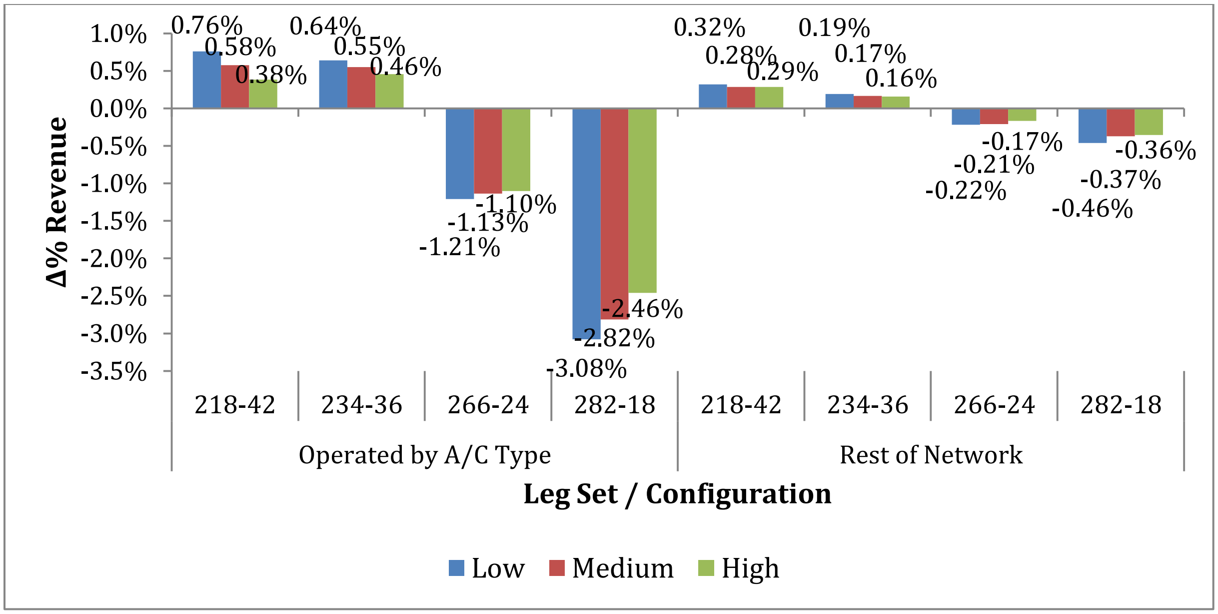

Figure 4-18: PODS - Proportional Revenue Variation by set of legs – Aircraft Type C – OD Control

The same pattern as for aircraft types A and B are observed for aircraft type C: Table 4-13 and Table 4-14 show that, in most cases, the BSM is able to determine correctly if a configuration change will be beneficial for an airline (or not). However, in the cases where the premium cabin capacity is increased at the high demand scenario for the OD-control baseline, it is observed that the BSM predicts loses while the simulation results in gains.

		Configuration			
Demand	**Mechanism**	**218-42**	**234-36**	**266-24**	**282-18**
Low	BSM	2.36%	1.51%	-2.19%	-4.81%
	PODS	0.88%	0.66%	-1.32%	-3.33%
Medium	BSM	1.94%	1.25%	-1.92%	-4.20%
	PODS	0.63%	0.56%	-1.20%	-2.91%
High	BSM	1.49%	0.97%	-1.60%	-3.52%
	PODS	0.41%	0.52%	-1.09%	-2.57%

Table 4-13: Proportional revenue variations on legs operated by aircraft C – Leg-based RM

		Configuration			
Demand	**Mechanism**	**218-42**	**234-36**	**266-24**	**282-18**
Low	BSM	1.35%	1.06%	-1.81%	-4.12%
	PODS	0.76%	0.64%	-1.21%	-3.08%
Medium	BSM	0.79%	0.77%	-1.51%	-3.56%
	PODS	0.58%	0.55%	-1.13%	-2.82%
High	BSM	-0.90%	-0.13%	-0.54%	-1.65%
	PODS	0.38%	0.46%	-1.10%	-2.46%

Table 4-14: Proportional revenue variations on legs operated by aircraft C – OD-control RM

4.4 Differences between Analytical Model and Simulations

Based on the results presented in sections 4.2 and 4.3, there are significant differences between the results obtained by the dual cabin BSM model when compared to the simulations. In general, the BSM model predicts adequately whether the change of the configuration of an aircraft type operating a subset of the legs will result in gains or losses for the airline on such specific legs. However, the proportional changes are typically exaggerated by the dual cabin BSM model; furthermore, the model completely ignores the effect of a configuration change on the rest of the network. This section aims to understand the main sources of differences between the model and the simulations by examining two dimensions: average fares used for valuating spilled and additional passengers and estimation of load factors.

4.4.1 Spilled/Additional Passenger Fares

Two cases are used as examples and examined in detail in this section: modifications to Aircraft Type A and modifications of Aircraft Type C when AL1 uses leg-control at the medium demand level. The differences between the shares of passengers across fares in the baseline of each of these cases provide insights for understanding the differences between the results of the simulations and the dual cabin BSM.

- **Aircraft Type A (150–20) – Leg-Control - Medium Demand**

As discussed earlier in this chapter, flight legs operated by aircraft type 'A' combine low ALLF's in premium cabin and high ALLF's in economy cabin in the baseline scenario. In a breakout of the loads on such flights by fare classes, it is observed that more than half of the economy cabin passengers are booked in FC10 (Figure 4-19). Furthermore, the number of passengers in the premium fare classes represents less than 6% of the total passengers for such legs. However, since the FC10 fares are low compared to the other fare classes, the revenue distribution by fare class has a different shape: still, FC10 represents the highest share of AL1's revenue (32%), but other fare classes (mainly FC5, FC6, FC9 and FC2) make significant contributions to the revenue as well (Figure 4-20).

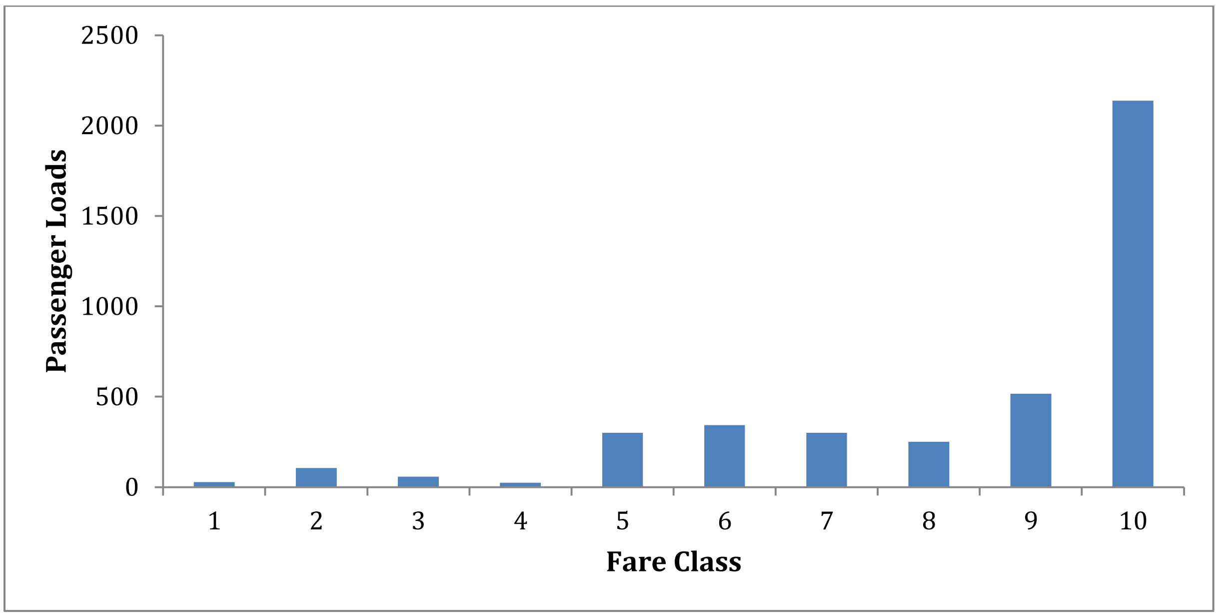

Figure 4-19: Passenger Loads by Fare Class in the Baseline – Legs Operated by Aircraft Type A - Leg Control - Medium Demand

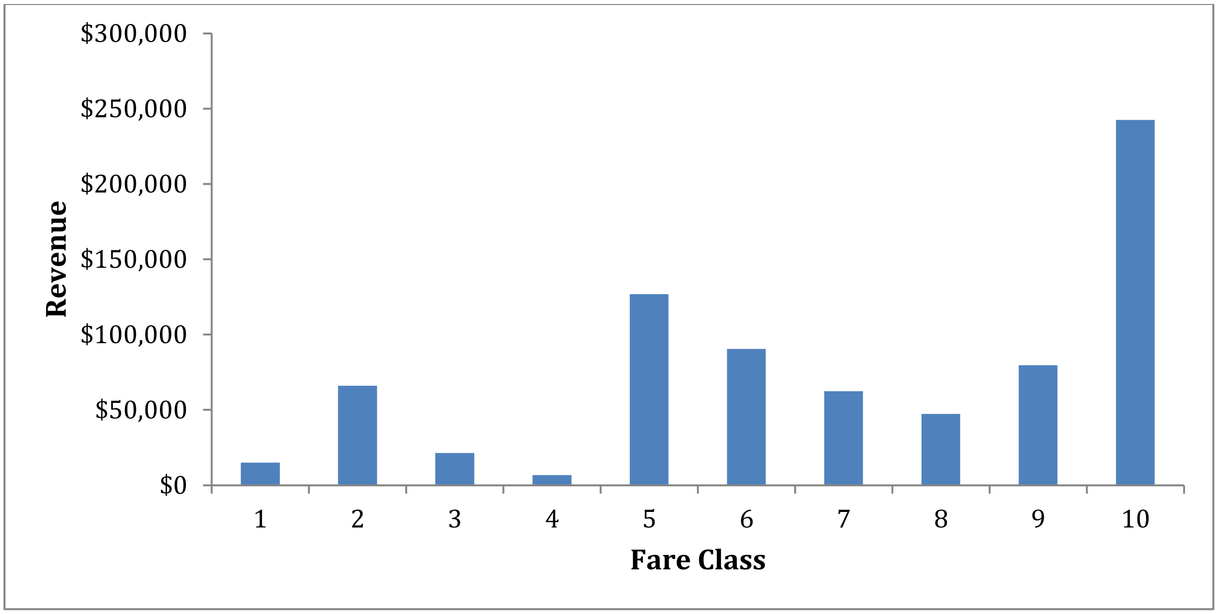

Figure 4-20: Revenue by Fare Class in the Baseline – Legs Operated by Aircraft Type A - Leg Control - Medium Demand

In order to support the comparison between the dual cabin BSM and the simulations, average fares for each of fare classes and cabin are calculated based on the information of revenue and passengers (Table 4-15). While the premium cabin average fare is between the average fares of FC2 and FC3, the economy cabin average fare is between FC8 and FC9.

Fare Class	Average Fare	Cabin	Cabin Average Fare
FC1	$531.37	Premium	$508.57
FC2	$627.55		
FC3	$373.08		
FC4	$280.86		
FC5	$423.00	Economy	$168.67
FC6	$264.41		
FC7	$207.98		
FC8	$188.18		
FC9	$154.16		
FC10	$113.39		

Table 4-15: Average Fare by Fare Class and Cabin in the Baseline - Legs Operated by Aircraft Type A – Leg Control – Medium Demand

As the configuration of aircraft type A is modified, the loads by fare classes also change. However, instead of distributing equally between fare classes, most of the variations in passenger loads are observed in FC10 (Figure 4-21). As a matter of fact, this is a result of the effective implementation of RM systems: on the one hand, passengers with low willingness to pay (who book in FC10) that would have been rejected in the baseline scenario are able to book only when additional capacity is made available; on the other hand, many of such FC10 passengers are rejected when economy cabin decreases.

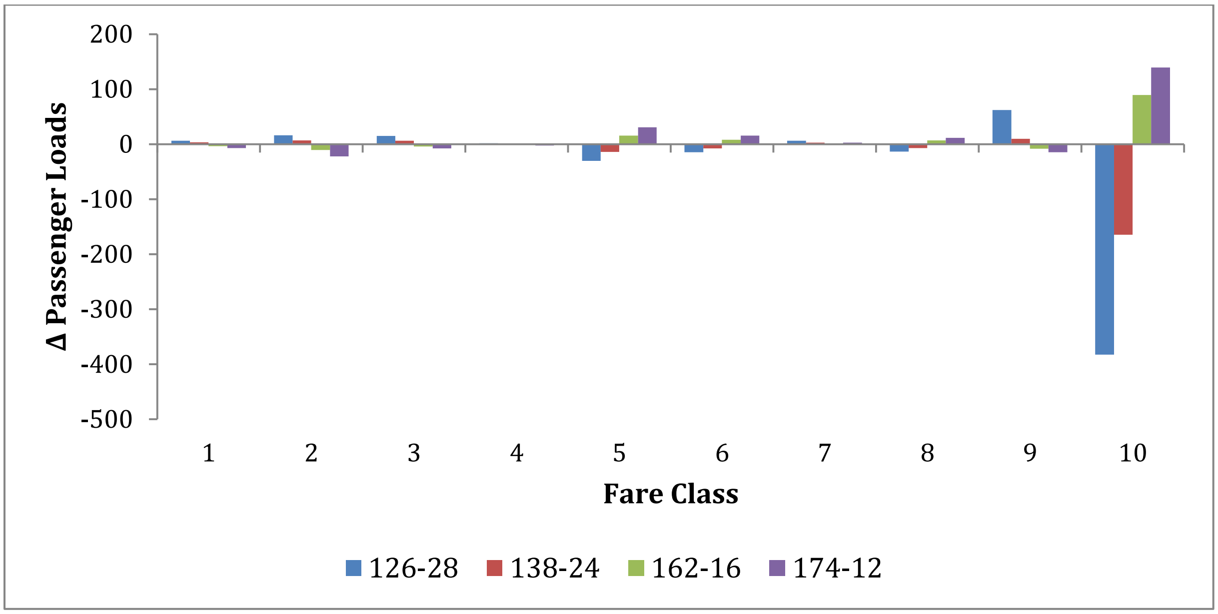

Figure 4-21: Absolute Variation of Passenger Loads by Fare Class – Legs Operated by Aircraft Type A - Leg Control - Medium Demand

When considering the prorated fare of the passengers accepted or rejected as a result of the configuration changes, it is found that FC10 is also the main source of variation of AL1's revenue (Figure 4-22). However, at least two other relevant findings related with the economy fare classes are presented: first, other fare classes (in this case FC9) increase revenue despite the reduction in economy cabin capacity. This suggests that some passengers that would have flown in FC10 but are willing to pay a higher fare are forced to book in FC9. Second, significant revenue losses are also observed in the highest economy fare class (FC5). This result is not expected, as RM systems in place are supposed to protect capacity for these fare classes; nevertheless, as "demand for premium fare classes" is being spilled because of the reduced capacity in premium cabin, some of those passengers are forced to buy in the less restricted economy fare class (others could decide as well to fly with a competitor or not to fly at all). In contrast, when premium cabin capacity is increased, passengers that would have booked in FC5 in the baseline configuration are more likely to find premium fare classes that provide a higher utility than FC5. This finding is one of the key differences with the dual cabin BSM, as the simulation allows identifying interactions between cabins that are not considered when cabins are modeled independently.

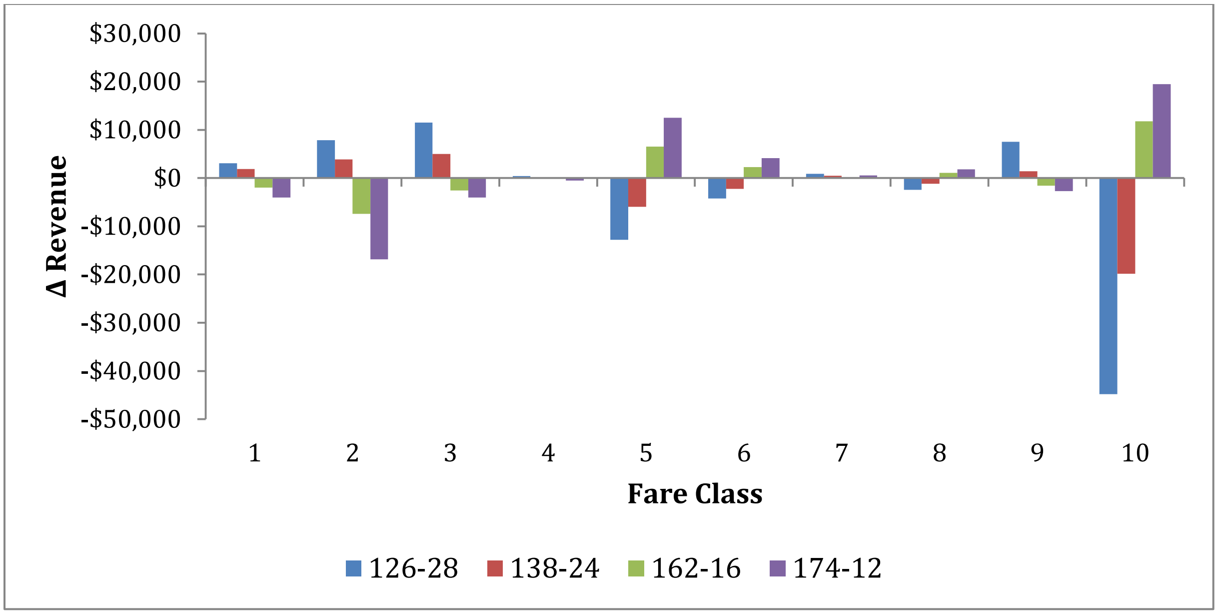

Figure 4-22: Absolute Variation of Revenue by Fare Class – Legs Operated by Aircraft Type A - Leg Control - Medium Demand

With respect to the premium fare classes (FC1 to 4) it is observed that FC2 and FC3 make a substantial proportion of the revenue gains when additional capacity is made available; by contrast, when premium cabin capacity is reduced, most of the revenue losses are observed in FC2. This difference illustrates another difference between the simulations and the dual cabin BSM: as changes are observed across all fare classes, it is unlikely that the average revenue captured from a passenger that is accommodated is the same as the average revenue lost from a passenger that is spilled because of a reduction in the cabin capacity.

Based on the changes in passenger loads and revenue by fare class and cabins from the PODS tests, the average revenue added (or spilled) by each additional passenger in each cabin can be calculated for each configuration with respect to the baseline configuration of 150 seats in economy cabin and 20 seats in premium cabin (Table 4-16, which includes also the results for the low and medium demand scenarios). When compared to the weighted fare (80% lowest. 20% average) used in the dual cabin BSM (within each cabin), it is found that the simulated average fare for each premium cabin passenger accommodated or spilled is much higher (around $600-$700 compared to weighted fare of roughly $320 in the BSM). It is also found that these values decrease as the demand increases. Although the differences are lower, in most of the cases the dual cabin BSM also underestimates the estimated impact on revenue of every passenger spilled or added in economy cabin because of a change in the configuration.

		Dual Cabin BSM	**PODS**			
		Weighted Fare	**Average Fare of Accommodated/Spilled Passengers**			
Cabin	**Demand**		**126-28**	**138-24**	**162-16**	**174-12**
Premium	**Low**	$324	$662	$760	$678	$684
	Medium	$326	$586	$642	$647	$642
	High	$320	$604	$669	$600	$579
Economy	**Low**	$126	$178	$181	$197	$218
	Medium	$124	$150	$151	$180	$193
	High	$124	$120	$121	$128	$141

Table 4-16: Average Revenue by passenger accommodated/spilled – Aircraft Type A – Leg-control

- **Aircraft Type C (250–30) – Leg-Control - Medium Demand Scenario**

Operating mostly international flights with high leg load factors in both premium and economy cabin, the legs operated by aircraft type C present a more uniform distribution passengers between the economy fare classes and a significant number of passengers in premium fare classes, mostly in FC2 (Figure 4-23); moreover, the main sources of revenue are FC2 and FC5 (Figure 4-24).

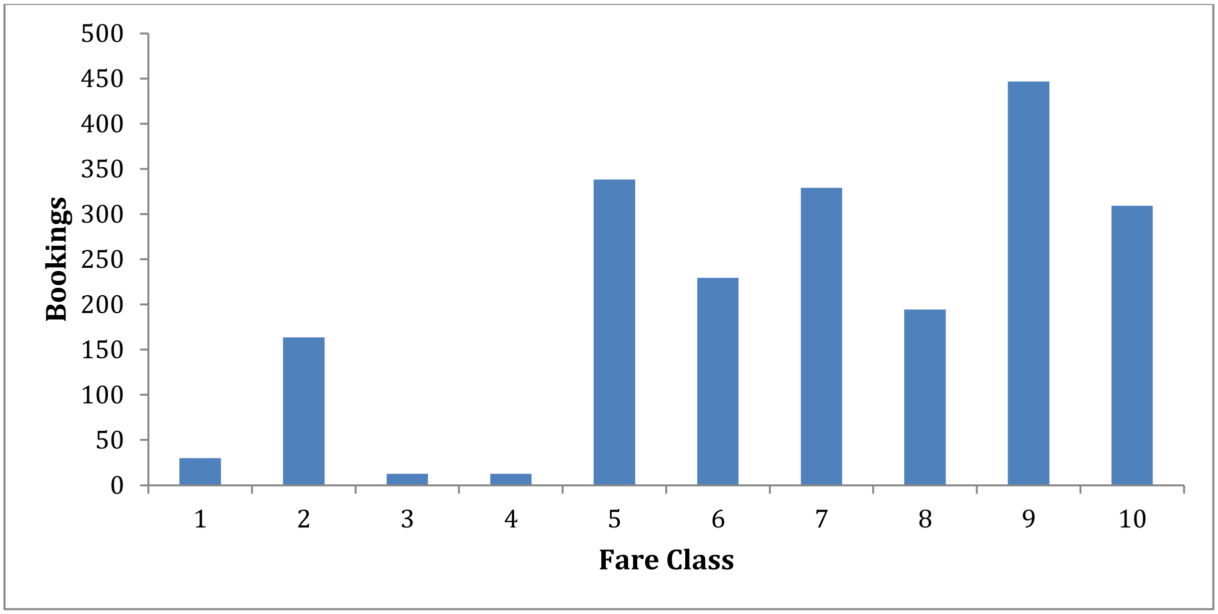

Figure 4-23: Passenger Loads by Fare Class in the Baseline – Legs Operated by Aircraft Type C - Leg Control - Medium Demand

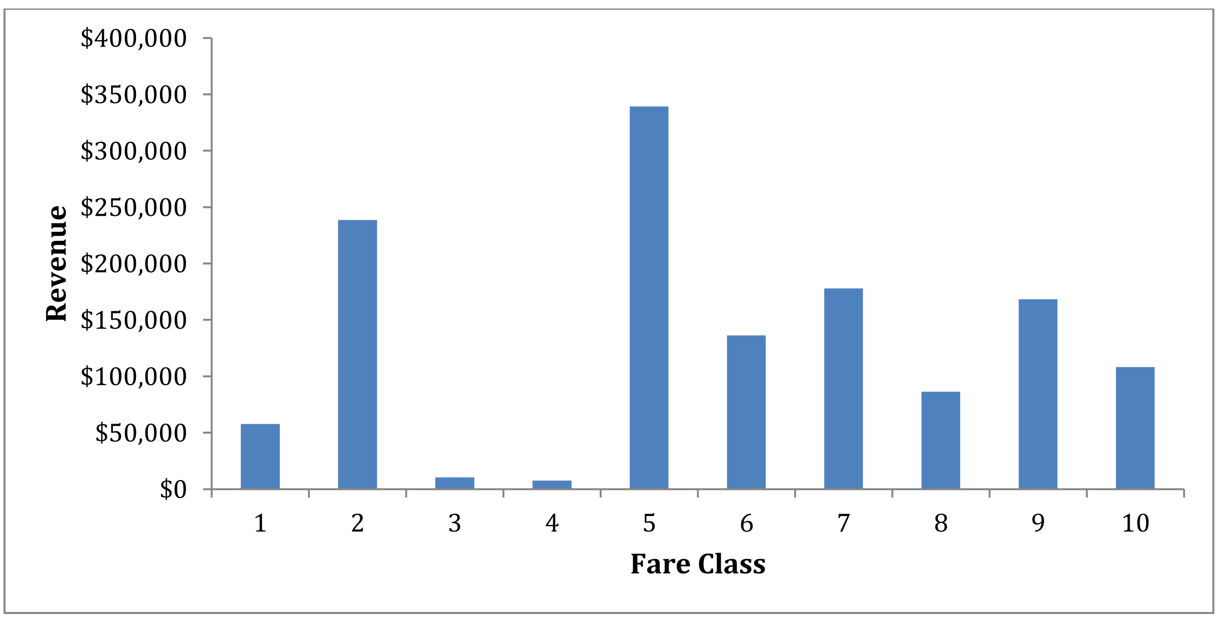

Figure 4-24: Revenue by Fare Class in the Baseline – Legs Operated by Aircraft Type C - Leg Control - Medium Demand

Table 4-17 shows that the average premium cabin fare is almost the same as the average FC2 and that the average economy cabin fare class is between the averages of FC6 and FC7.

Fare Class	Average Fare	Cabin	Cabin Average Fare
FC1	$1,924.68	Premium	$1,438.62
FC2	$1,459.82		
FC3	$832.16		
FC4	$600.80		
FC5	$1,002.19	Economy	$549.97
FC6	$594.02		
FC7	$539.98		
FC8	$443.93		
FC9	$377.01		
FC10	$349.51		

Table 4-17: Average Fare by Fare Class and Cabin in the Baseline - Legs Operated by Aircraft Type C – Leg Control – Medium Demand

While the variation in passengers in the aircraft type A case was mostly concentrated on FC10, in the case of aircraft type C substantial changes are also observed in other fare classes, such as FC2, FC3, FC5 and FC9 (Figure 4-25). However, considering that the average fare for FC10 is considerably lower than for the other fare classes, the main variations in terms of revenue are in FC2 and FC5. Again, this is different from the expectation, because of RM systems, of observing most of the variations in the lowest fare classes (as is the assumption in the dual cabin BSM). Instead, these results confirm that the modification of the capacity of the cabins leads to exchanges of passengers between the cabins, and many of these changes happen at the higher fare classes.

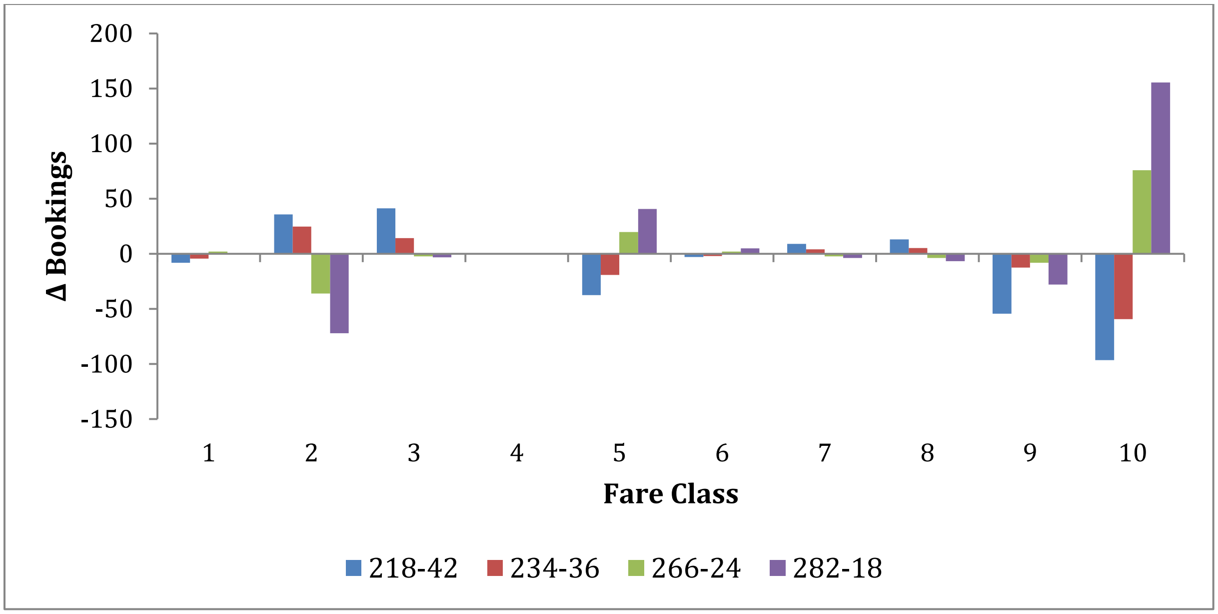

Figure 4-25: Absolute Variation of Passenger Loads by Fare Class – Legs Operated by Aircraft Type C - Leg Control - Medium Demand

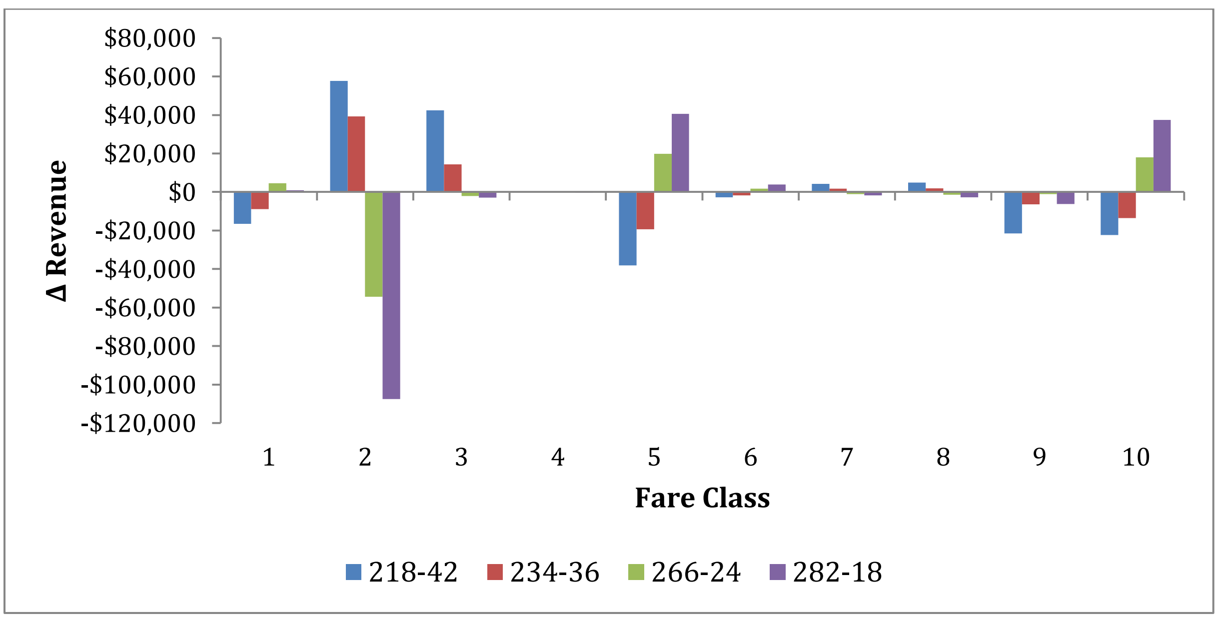

Figure 4-26: Absolute Variation of Revenue by Fare Class – Legs Operated by Aircraft Type C - Leg Control - Medium Demand

In summary, the weighted fare used in the dual cabin BSM is significantly lower than the average fare of each accommodated/spilled passenger found in the simulations, as presented in Table 4-18. This is true for both cabins, but the difference is larger for the premium cabin.

		Dual Cabin BSM	PODS			
		Weighted Fare	Average Fare of Accommodated/Spilled Passengers			
Cabin	**Demand**		**218-42**	**234-36**	**266-24**	**282-18**
Premium	**Low**	$998	$1,137	$1,170	$1,339	$1,394
	Medium	$953	$1,216	$1,285	$1,425	$1,443
	High	$953	$1,300	$1,388	$1,431	$1,441
Economy	**Low**	$356	$425	$433	$507	$561
	Medium	$390	$447	$444	$434	$439
	High	$438	$447	$454	$458	$460

Table 4-18: Average Revenue by passenger accommodated/spilled – Aircraft Type C – Leg-control

- **Other scenarios**

The results for the other scenarios are reported in Table 4-19, Table 4-20, Table 4-21 and Table 4-22. As in the scenarios reported above, the simulation average fare for each accommodated or spilled premium cabin passenger is always higher than the weighted fare used in the dual cabin BSM; as a matter of fact in some scenarios the value is higher than the average fare in the baseline. The results for economy cabin fares are diverse: in most of the cases the average fare of each accommodated/spilled passenger is between the average and the weighted fare, but there are also cases where the values are higher than the average fare and other cases where the weighted fare is higher than the average fare provided by the simulation.

			Dual Cabin BSM	PODS			
			Weighted Fare	Average Fare of Accommodated/Spilled Passengers			
Cabin	**Demand**	**Avg. Fare**		**126-28**	**138-24**	**162-16**	**174-12**
Premium	**Low**	$1,313	$994	$1,589	$1,681	$1,702	$1,603
	Medium	$1,301	$976	$1,580	$1,572	$1,552	$1,500
	High	$1,262	$895	$1,465	$1,438	$1,383	$1,388
Economy	**Low**	$415	$228	$315	$352	$439	$469
	Medium	$413	$222	$298	$312	$351	$385
	High	$411	$218	$260	$271	$285	$304

Table 4-19: Average Revenue by passenger accommodated/spilled – Aircraft Type B – Leg-control

			Dual Cabin BSM	PODS			
			Weighted Fare	Average Fare of Accommodated/Spilled Passengers			
Cabin	**Demand**	**Avg. Fare**		**126-28**	**138-24**	**162-16**	**174-12**
Premium	**Low**	$515	$326	$636	$731	$691	$667
	Medium	$510	$319	$563	$633	$590	$590
	High	$504	$310	$504	$588	$534	$550
Economy	**Low**	$171	$127	$101	$105	$133	$145
	Medium	$174	$132	$103	$102	$104	$112
	High	$180	$140	$110	$108	$100	$106

Table 4-20: Average Revenue by passenger accommodated/spilled – Aircraft Type A – OD-control

			Dual Cabin BSM	PODS			
			Weighted Fare	Average Fare of Accommodated/Spilled Passengers			
Cabin	**Demand**	**Avg. Fare**		**126-28**	**138-24**	**162-16**	**174-12**
Premium	**Low**	$1,316	$990	$1,469	$1,496	$1,503	$1,483
	Medium	$1,311	$971	$1,497	$1,503	$1,435	$1,427
	High	$1,296	$904	$1,376	$1,400	$1,367	$1,376
Economy	**Low**	$408	$232	$277	$292	$409	$458
	Medium	$410	$238	$282	$291	$347	$372
	High	$416	$256	$286	$293	$330	$336

Table 4-21: Average Revenue by passenger accommodated/spilled – Aircraft Type B – OD-control

			Dual Cabin BSM	PODS			
			Weighted Fare	Average Fare of Accommodated/Spilled Passengers			
Cabin	**Demand**	**Avg. Fare**		**126-28**	**138-24**	**162-16**	**174-12**
Premium	**Low**	$1,496	$809	$1,124	$1,191	$1,317	$1,395
	Medium	$1,520	$803	$1,204	$1,258	$1,390	$1,444
	High	$1,535	$802	$1,274	$1,322	$1,469	$1,471
Economy	**Low**	$529	$356	$510	$540	$592	$638
	Medium	$537	$368	$454	$454	$473	$497
	High	$549	$389	$439	$444	$457	$474

Table 4-22: Average Revenue by passenger accommodated/spilled – Aircraft Type C – OD-control

4.4.2 Load Factor Estimation

The estimation of leg load factors represents another source of differences between the dual cabin BSM model and the results of the simulations. Leg load factors are calculated in BSM using the methodology described in Chapter 3 and briefly discussed in section 4.2. Those results are compared in this section with the output of the simulations and analyzing in detail the cases of Aircrat Type A and Aircraft Type C. The operation of Aircraft Type B is a mix between the domestic operation of A and the international operation of C and therefore it is not examined in detail in this section.

- **Aircraft Type A (Baseline Configuration: 150–20) – Leg-Control**

Figure 4-27 presents the relationship between premium and economy leg load factors when the capacity of aircraft type A is modified by increasing the premium cabin seat capacity and reducing the economy cabin seat capacity, to the 126-28 configuration. The diagonal line in the figure shows the point where the leg load factor calculated using the dual cabin BSM model would be the same as the leg load factors observed in the PODS simulations. Points below the diagonal identify legs for which the BSM model estimated higher load factors than the ones observed in the PODS simulations, while points above the diagonal indicate that the average load for a leg in the simulations is found to be higher than predicted by the BSM.

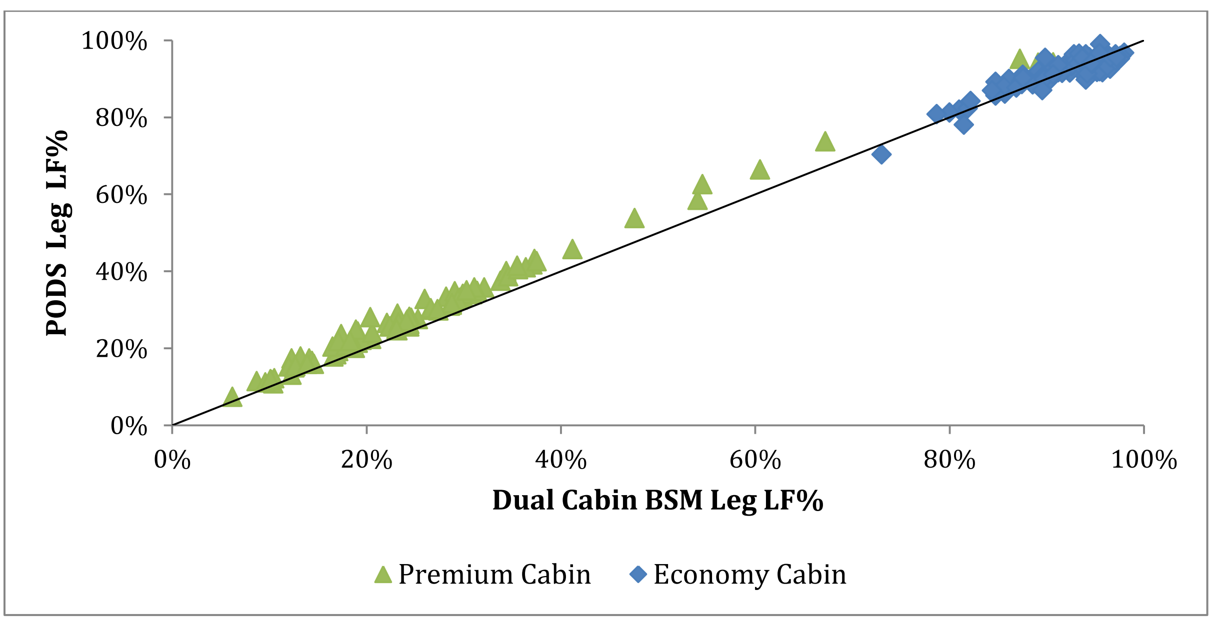

Figure 4-27: Leg Load Factor Estimation for aircraft type A in Configuration 126-28 – Leg Control

Following this approach, it is observed that most of the premium cabin leg load factors are underestimated by the dual cabin BSM. In a similar way, the graphs suggest that the BSM underestimates also the economy cabin leg load factors in most of the cases; however, this is less clear at very high load factors (over 90%), where the points seem to be distributed equally above and below the diagonal. The percentage points difference between the BSM and PODS at different levels of load factor are also detailed in Figure 4-28. In this case, negative load factor differences (measured in percentage points) indicate that the BSM calculates lower values than what was found in the simulation.

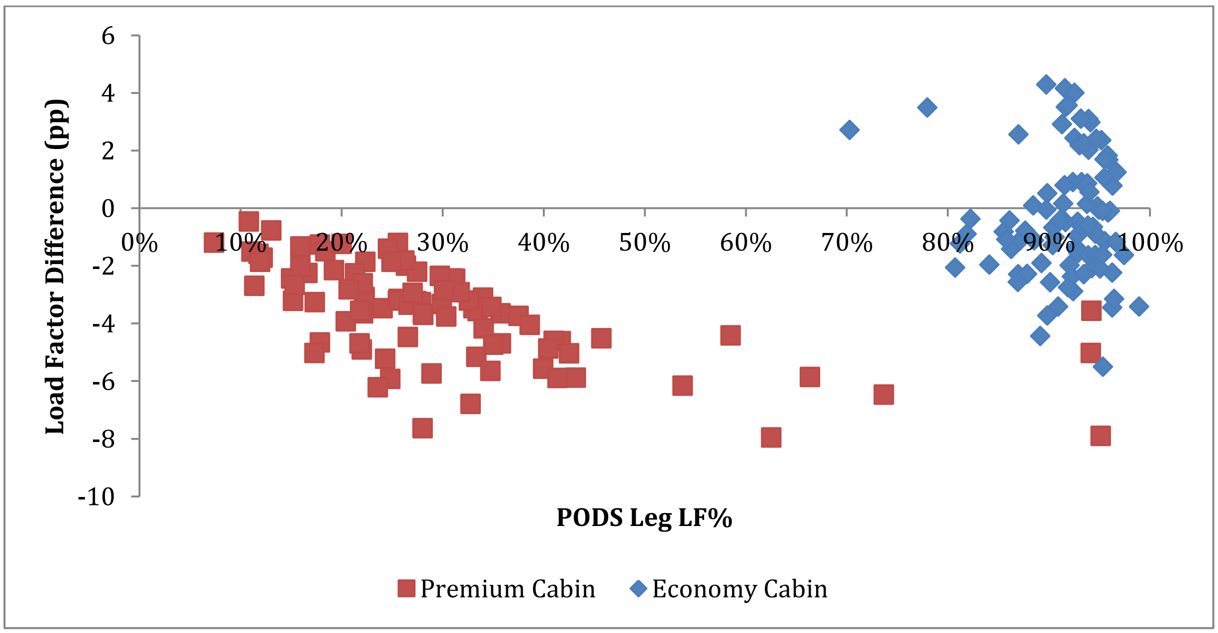

Figure 4-28: Load Factor differences between BSM and PODS - Aircraft type A in Configuration 138-24 – Leg-based RM

Leg load factors obtained in PODS are also compared to the ones found by the dual cabin BSM for the 174-12 configuration and are illustrated in Figure 4-29. In this case, the BSM tends to overestimate the load factor of both cabins of each leg when compared to the outcome of the simulations. While this trend is particularly clear in the case of premium cabin, the finding is less clear on the legs with economy cabin load factors that exceed 90%.

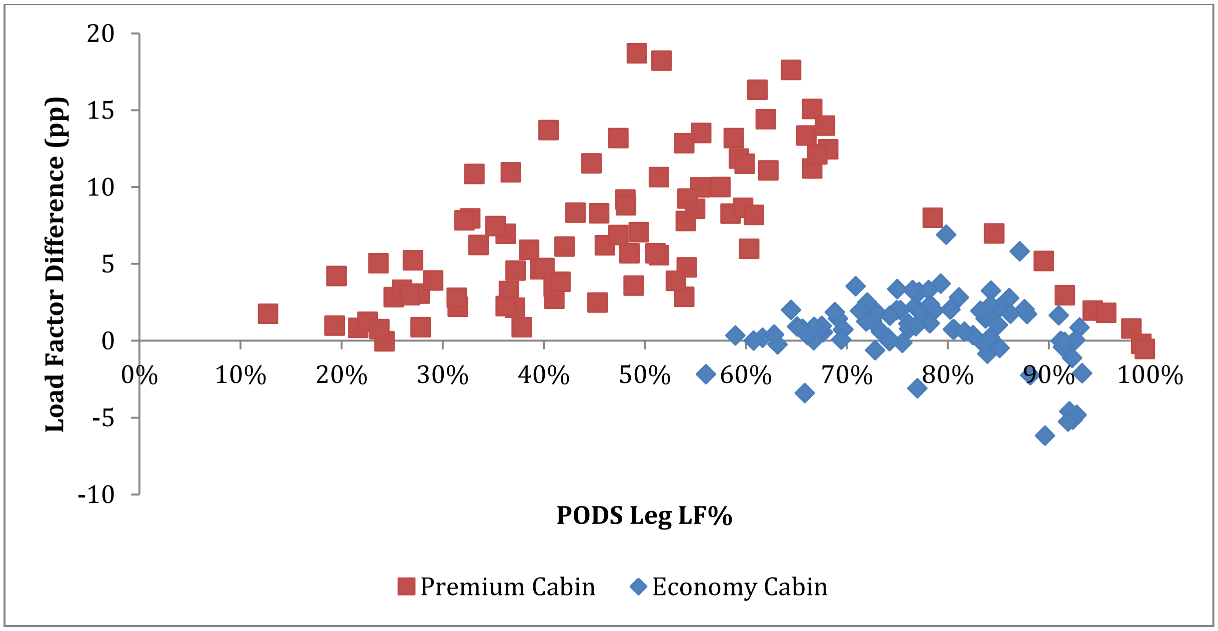

Figure 4-29: Leg Load Factor Estimation for aircraft type A in Configuration 174-12 – Leg Control

- **Aircraft Type C (Baseline Configuration: 250–30) – Leg-Control**

Following the same logic discussed in the previous section, Figure 4-30 compares the load factors found in the simulations against the ones calculated using the dual cabin BSM for the modifications that increase the size of the premium cabin (and reduce the size of economy cabin) of aircraft type C. The results for the premium load factors are similar to the ones found for aircraft type A as the dual cabin BSM tends to underestimate their values. However, a different behavior is observed for economy cabin load factors: while in the case of aircraft type A the BSM tended to underestimate them, in the case of aircraft type C it leans towards overestimation when economy cabin capacity is reduced.

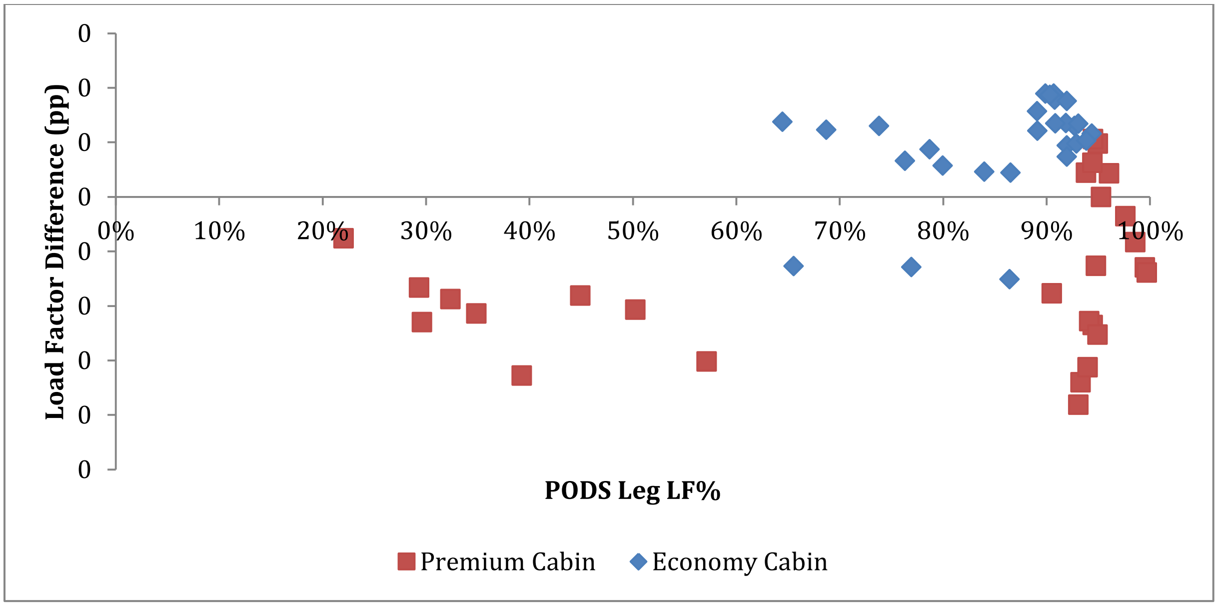

Figure 4-30: Load Factor differences between BSM and PODS - Aircraft type C in Configuration 218-42 – Leg Control

The opposite happens when the economy cabin capacity is increased with respect to the baseline (Figure 4-31). Premium cabin load factors are overestimated by the dual cabin BSM (as was the case for aircraft type A) while the economy cabin load factors are underestimated (different from what was observed for aircraft type A).

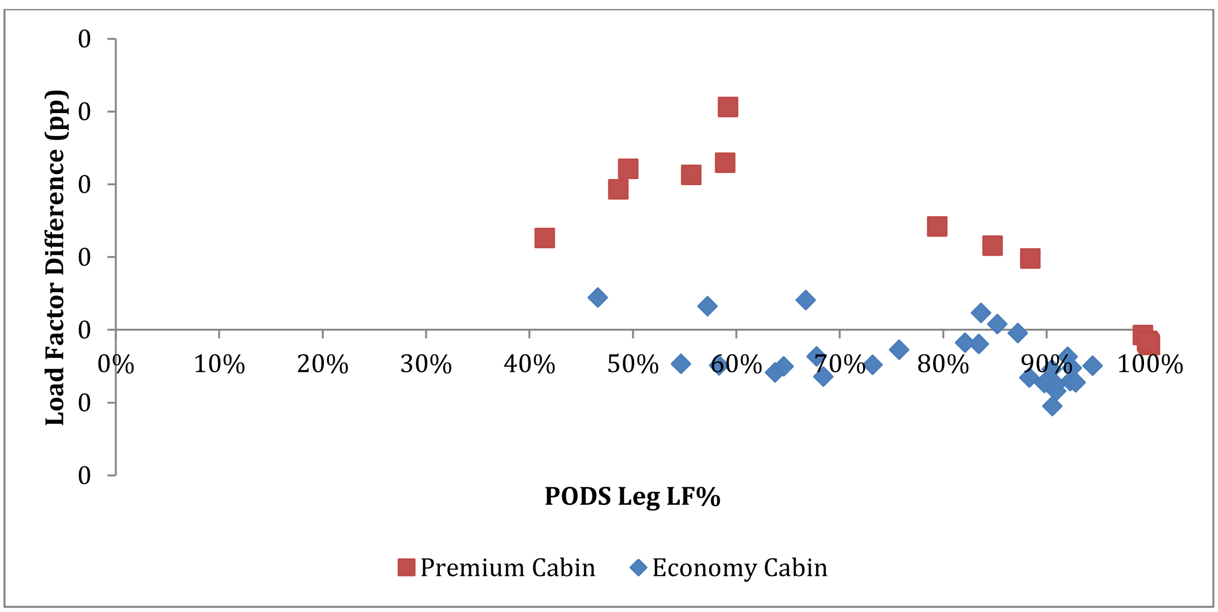

Figure 4-31: Load Factor differences between BSM and PODS - Aircraft type A in Configuration 282-18 – Leg Control

In order to provide an objective measurement of bias that compares the leg load factors estimated using the dual cabin BSM and the results in PODS across aircraft types, configurations, RM methods and demand levels, a simple measure of error is proposed: the average of the difference in percentage points, $\text{ALFD}_{\mathcal{L}}$, compares the load factor estimated by the BSM on leg i for cabin j, $BSMLF_{i,j}$, and the load factor found in the simulation for the same cabin and leg, $PODSLF_{i,j}$, for a set of legs $\mathcal{L}$ and a the set of cabins $\mathcal{C} = \{premium, economy\}$.

$$LFDF_{i,j} = BSMLF_{i,j} - PODSLF_{i,j}$$

$$ALFD_{\mathcal{L}} = \sum_{i\in\mathcal{L}, j\in\mathcal{C}} LFDF_{i,j}/|\mathcal{L}| \times 100$$

The values for $ALFD$ for each scenario modeled for aircraft types A, B and C are presented in Figure 4-32, Figure 4-33 and Figure 4-34, respectively. These graphs help to identify some relevant patterns: first, the $ALFD$'s for premium cabin load factors are typically higher than the $ALFD$'s for economy cabin load factors. Second, the premium cabin $ALFD$'s are always positive when premium cabin capacity decreases and negative when premium cabin capacity increases; therefore, the BSM model tends to underestimate the premium cabin load factors when the premium cabin capacities increase and to overestimate the premium cabin load factors when capacities decrease. Similarly, in most of the cases it is found that the BSM also underestimates the economy cabin load factors when the number of seats in economy cabin decreases and overestimates it when the capacity increases.

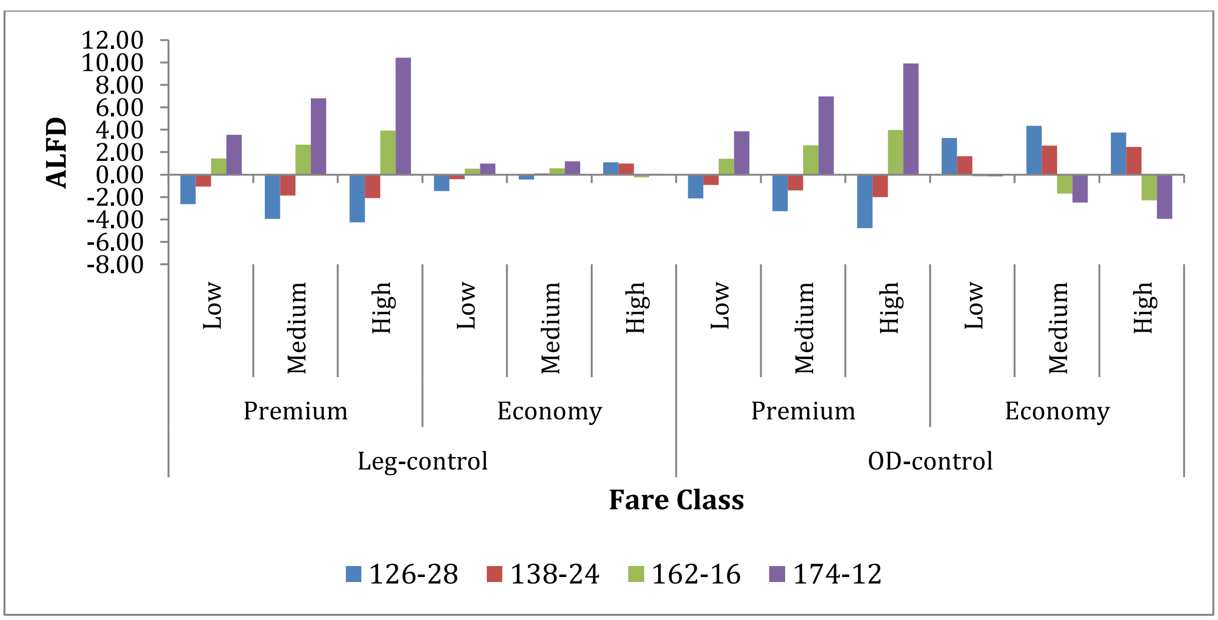

Figure 4-32: Average of the difference in percentage points – Aircraft Type A

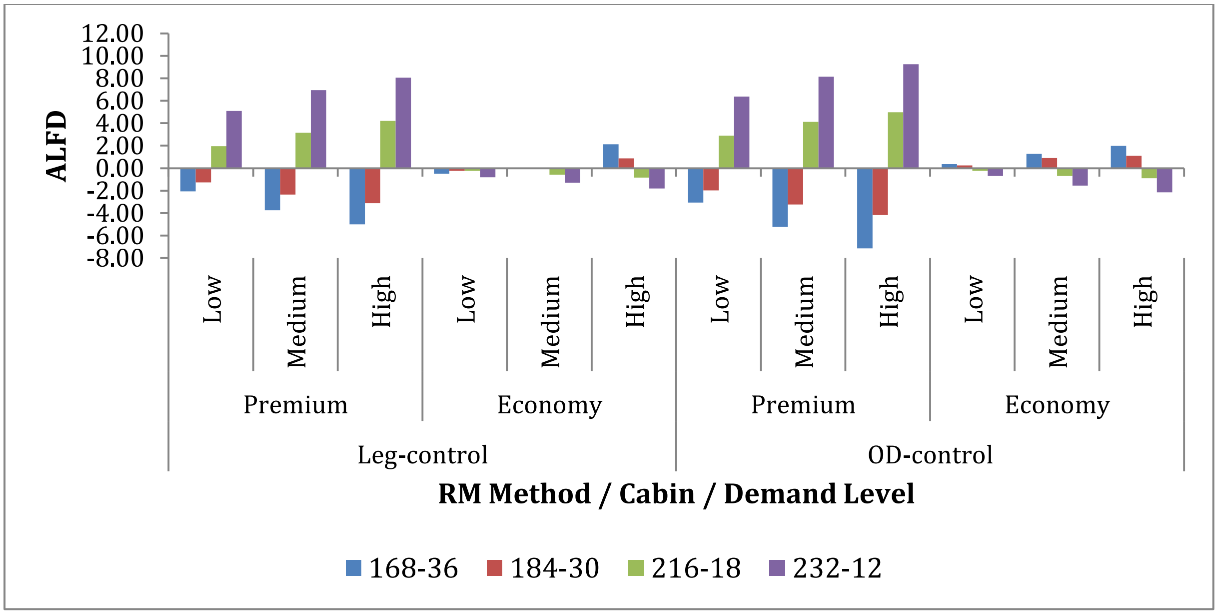

Figure 4-33: Average of the difference in percentage points – Aircraft Type B

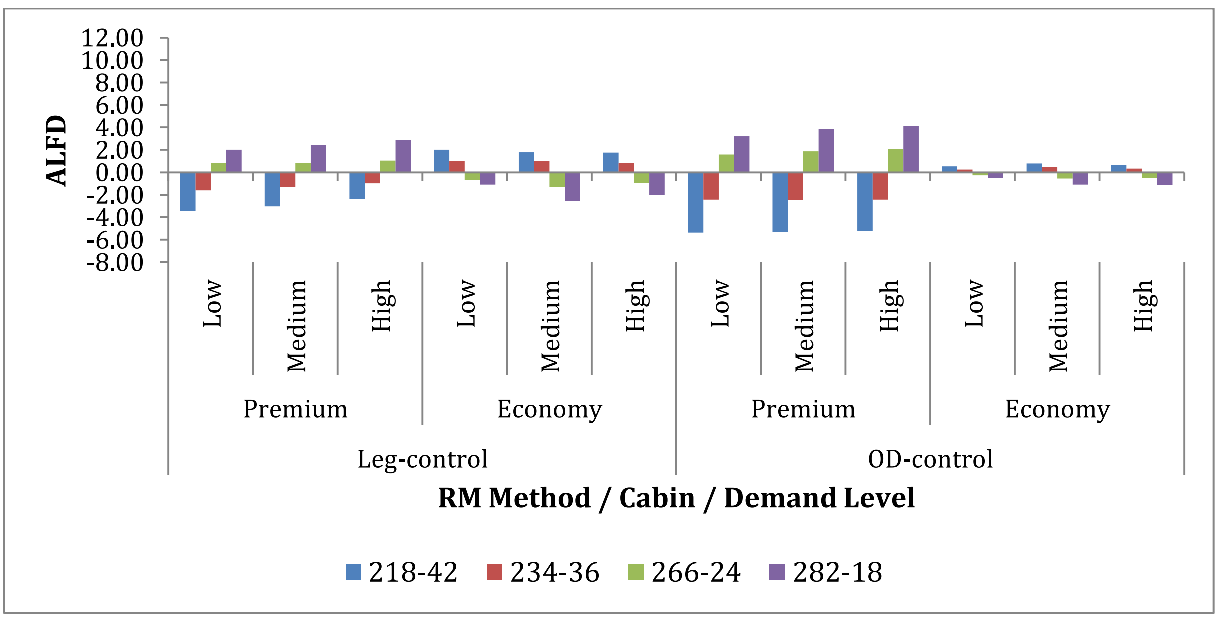

Figure 4-34: Average of the difference in percentage points – Aircraft Type C

As with the fares, the interaction between the sizes of the cabins offers, at least, a partial explanation for the errors of the BSM determining the load factor obtained from the simulations. For example, when the economy cabin size of aircraft type C is increased from 250 to 266, the BSM model would determine the number of passengers to be carried in the new configuration by estimating the unconstrained demand for economy cabin and adjusting it to the modified capacity. However, the model does not consider that, simultaneously, the capacity of the premium cabin is reduced from 30 to 24, spilling premium cabin passengers. Some of those passengers would consider flying in economy cabin, adding demand to the unconstrained demand for economy cabin seats contemplated initially by the model (explaining why the model underestimates the increases when capacity increases). In addition to the passengers spilled from premium cabin, some premium cabin passengers would have the option to book in economy cabin thanks to the additional capacity made available. Hence, the unconstrained demand for premium cabin would also be reduced because of such passengers that would prefer to book in an economy fare class if available (explaining why the model overestimates when capacity is reduced). Of course, the same logic would apply if the premium cabin capacity was increased while economy cabin capacity decreased.

This pattern is not observed for economy cabin of aircraft types A and B, in the case of low and medium demand levels with leg control RM. As discussed before, the premium cabin ALLF's of those legs is low (<50%). Considering these low load factors, it could be said that very few of the passengers flying in economy cabin were forced to fly in economy cabin because of lack of availability in premium cabin; therefore, if additional premium cabin capacity is made available it is unlikely to observe the additional demand for premium cabin fare classes described in the previous paragraph. Hence, the model is less likely to underestimate the load factor of economy cabin. In general, the load factor of a cabin should be moderate or high (>50%) in order to generate additional demand to the other cabin when the capacity of such other cabin is increased.

4.5 Summary

This chapter has examined the impacts on the revenues that result from the modification of the capacities of premium and economy cabins of an airline operating dual cabin aircraft. The analysis was initially conducted by applying an adaptation to the Boeing-Swan spill model for dual cabin aircraft, assuming that the demand for travel in one cabin is independent from the demand for travel in the other cabin. These results were then compared to the output of simulations performed in PODS, testing on three aircraft types with five alternate configurations for each one, at different demand levels and assuming scenarios where the airline was performing leg and OD-control in its revenue management systems.

The initial finding is that the dual cabin BSM is able to predict reasonably well whether a change of the configuration of an aircraft will generate revenue gains or losses for an airline. However, when the results of the simulation are examined specifically for the legs operated by the modified aircraft type, it is found that the dual cabin BSM tends to exaggerate both the revenue gains and the losses on such legs. In addition, it is found that the revenue changes in the rest of the network are substantial in some scenarios, but the dual cabin BSM does not take this into consideration.

In order to understand the differences between the analytical method and the simulations, the results of such simulations were examined in detail focusing mostly in two critical aspects: the average fare considered for each passenger that is either accommodated or spilled from each cabin because of the change of configuration and the estimation of the load factors. We found that changes in the capacity of a cabin do not affect only the bookings in the lowest fares classes applicable for each cabin, as would be expected when RM systems are in place. Instead, we found that there is an interaction between a cabin that increases its capacity and another one that reduces its number of seats causes an exchange of some passengers between cabins, affecting the bookings in higher fare classes as well. This finding is relevant as it suggests that the assumption of independence of the demand for each cabin is likely incorrect and that the assumption of a weighted fare (20% average fare - 80% lowest fare) that values the impact on the revenue of each passenger spilled or accommodated because of the additional capacity typically underestimates the effect on total revenue of the modification of the configuration.

The estimations of the load factors made by using the dual cabin BSM were also compared with the simulations in PODS. The main findings are that when the capacity of a cabin increases, the BSM typically underestimates the new load factors for such legs; likewise, when the capacity of a cabin is reduced, the BSM typically overestimates the expected new load factors. This finding is also associated with the relationship between the cabins, as the assumption of independent unconstrained demand for each cabin considers neither that passengers spilled from the other cabin represent additional demand nor the change of the decision of the passengers because of the changes in availability in each cabin.

5. Premium Cabin Capacity Sharing Schemes

In this chapter the results of the shared capacity methods introduced in Section 3.5 are presented and discussed. While the first section explores the impact of the implementation of such schemes when leg-based RM is used, the second section focuses on OD-control RM.

5.1 Leg-based RM Premium Cabin Capacity Sharing Schemes

This section explores Full EMSR, a strategy that is developed based on the logic of the leg-control RM optimizer EMSRb; this heuristic is described in detail in Section 3.5.1. The main characteristics of the baseline scenario for evaluating the performance of Full EMSR are described in Table 5-1. On the one hand, AL1 and AL3 use EMSRb as their technique for performing leg-based RM; on the other hand, AL2 and AL4 use ProBP as their mechanism for exercising OD-control RM. All the airlines compete in Network V1 and, when applicable, manage their cabins (premium and economy) distinctly, optimizing the revenue capture of each of these cabins independently. Since the aircraft configurations of AL3's fleet do not include any dual cabin aircraft, the concept of distinct or shared capacity does not apply for that airline. This baseline is simulated in PODS at three different demand levels: low, medium and high, as described in Chapter 3.

Leg-Control Base Case
Network V1
AL1: Distinct EMSRb - Leg-based Std. Forecasting
AL2: Distinct ProBP – Path-based Std. Forecasting
AL3: EMSRb – Leg- Based Std. Forecasting
AL4: Distinct ProBP – Path-based Std. Forecasting

Table 5-1: EMSRb Baseline Settings

5.1.1 <u>Full EMSR</u>

In order to evaluate the performance of Full EMSR, simulations are run with the scenario described in Table 5-2; 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 identifying the effects caused exclusively by the decision of AL1 to share the premium cabin capacity.

Leg-Control Full EMSR
Network V1
AL1: Full EMSR - Leg-based Std. Forecasting
AL2: Distinct ProBP – Path-based Std. Forecasting
AL3: EMSRb – Leg- Based Std. Forecasting
AL4: Distinct ProBP – Path-based Std. Forecasting

Table 5-2: Full EMSR Settings

The total revenue (premium + economy) proportional variations for each of the airlines in the market are presented in Figure 5-1. For Airline 1, this decision either has a negligible positive effect, as in the low demand scenario (increase of 0.03%), or is substantially negative and generate losses of up to -0.66%, as in the high demand scenario. It should be noted as well that none of the other airlines in the network is able to take advantage of the losses of AL1; as a matter of fact, all the airlines experience proportional revenue losses because of the decision of AL1 to share its premium cabin capacity with economy fare classes.

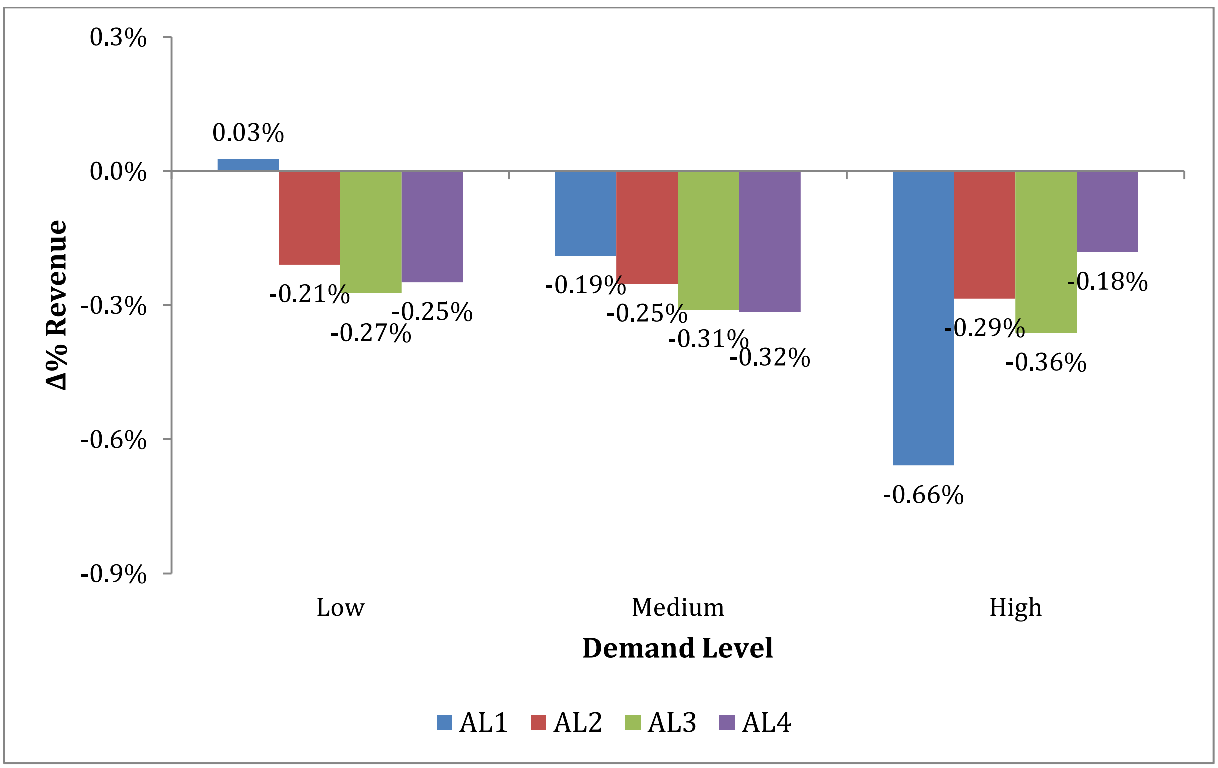

Figure 5-1: Full EMSR- Total Revenue Proportional Variation by Airline

The objective of implementing Full EMSR is to share the remaining capacity in premium cabin with additional passengers of economy fare classes; therefore, the expectation is to observe an increase in total revenue resulting from the growth of the revenue captured from the economy fare classes while keeping the same revenue from premium fare classes. However, when changes are presented by fare class type (Figure 5-2), it is observed that sharing premium cabin capacity with economy fare classes results in important proportional losses in the premium fare classes FC 1 to 4 (-5.81% to -14.59%), partially compensated by proportional gains for fare classes 5 to 10, ranging from 1.31% to 2.54%, depending on the demand level.

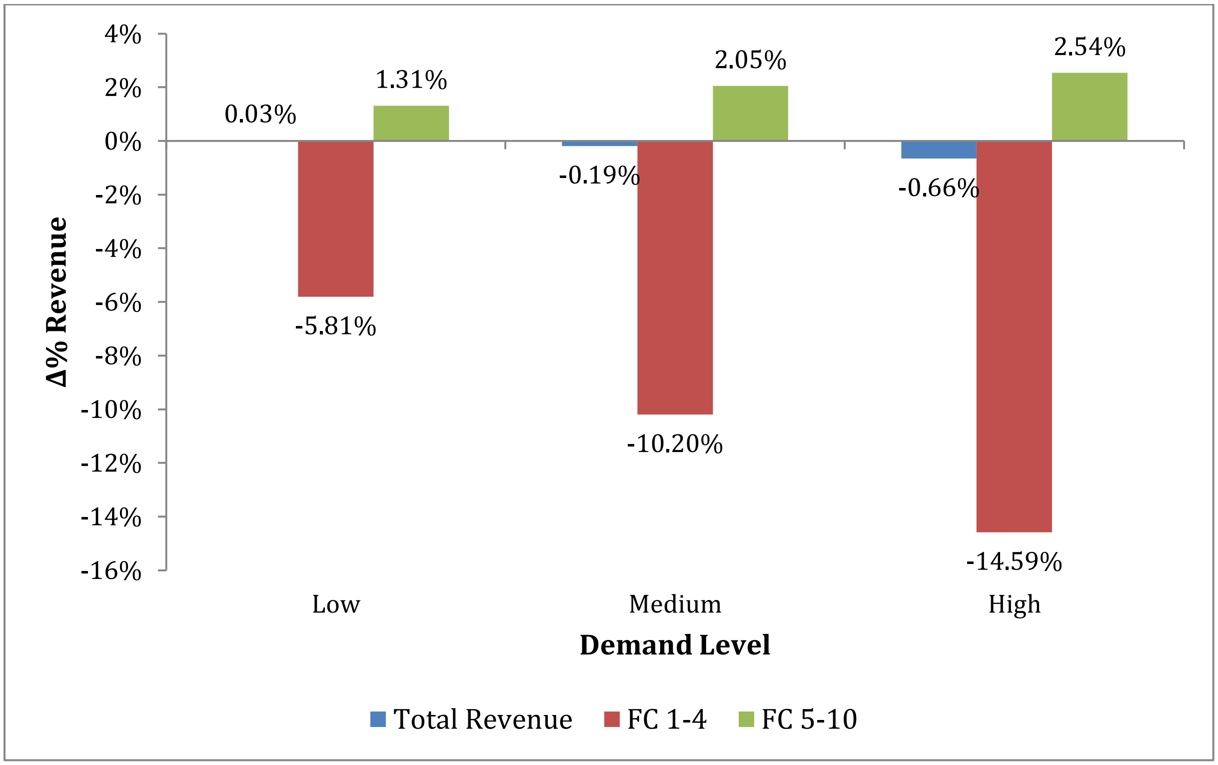

Figure 5-2: Full EMSR - Proportional Revenue Gains by Fare Class Type

Considering that the revenue from economy fare classes represents a large portion of AL1's total revenue (over 81%), the increase observed in the low demand scenario is just enough to compensate for the losses in premium fare classes; however, that is not the case for the medium and high demand scenarios, and that explains the total revenue losses observed in such cases (Figure 5-3).

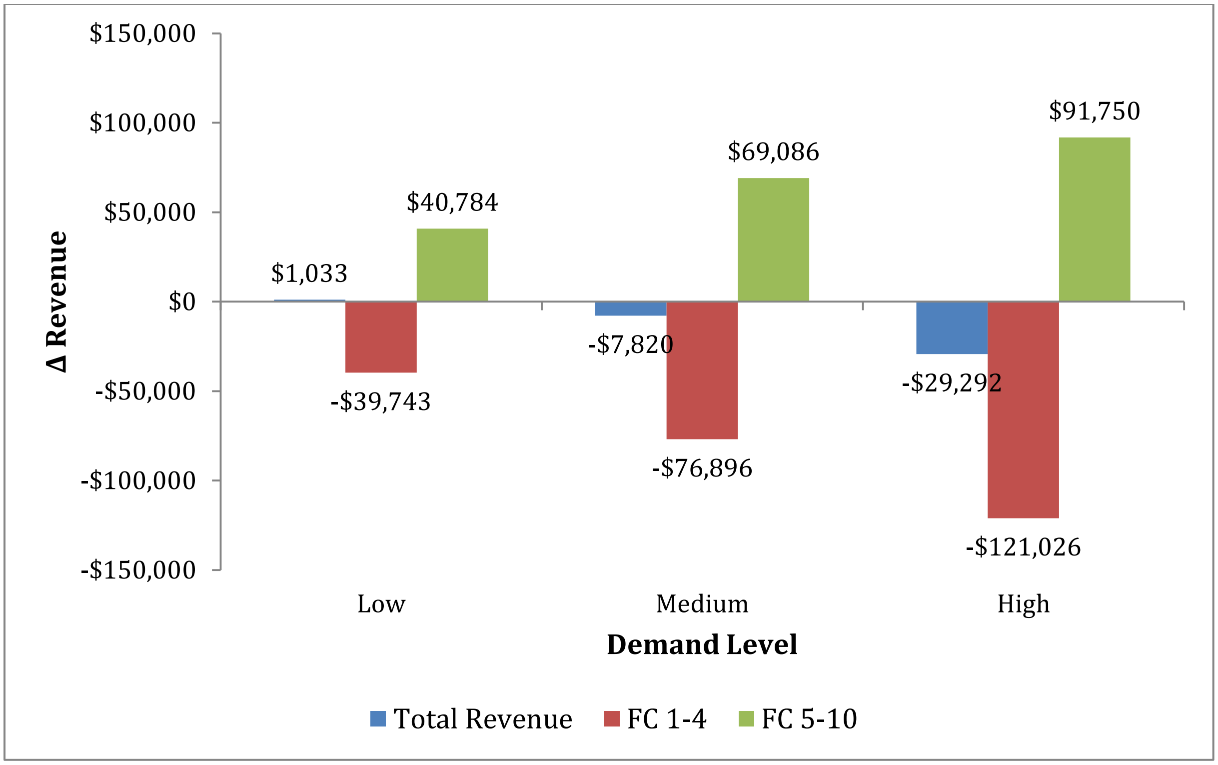

Figure 5-3: Full EMSR - Absolute Revenue Variations by Fare Class Type

Analyzing the absolute revenue variations by fare class type and cabin, as provided by Figure 5-4, is helpful to understand to the performance of Full EMSR. On the one hand, it achieves the objective of increasing the revenue captured from passengers flying in premium cabin by accommodating many passengers that book in economy fare classes; those increases compensate for the significant losses in revenue from passengers that booked in premium fare classes. On the other hand, it is found that the revenue from passengers that effectively fly in economy cabin decreases substantially. Table 5-3 summarizes the results of the simulations using Full EMSR and compares them with the results that airlines expect to achieve when a strategy of premium cabin capacity sharing is implemented; as presented in the table, the outcome differs substantially from the expectations.

The revenue variations by fare class are then reviewed in order to identify the sources of variation resulting from the use of Full EMSR. Based on Figure 5-5, FC1 and FC2 are the fare classes that suffer the higher revenue losses, while the revenue gains come mainly from FC5 and FC10.

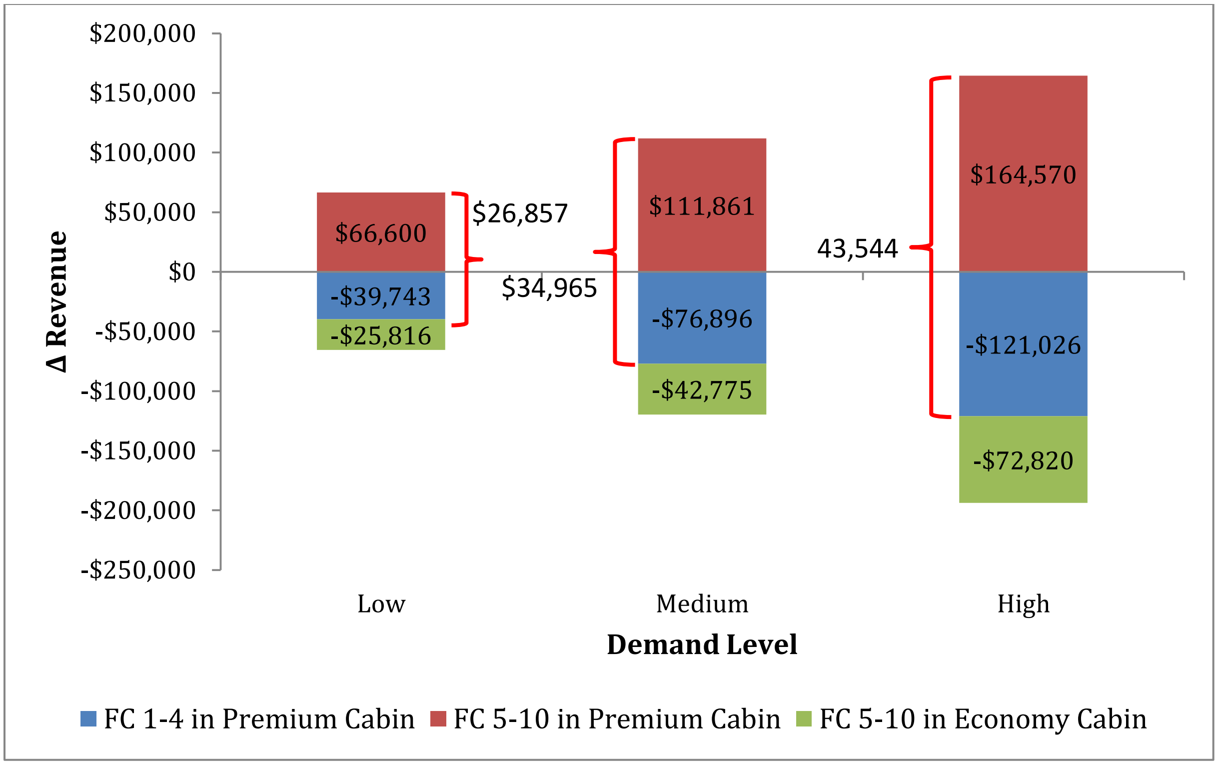

Figure 5-4: Full EMSR - Absolute Revenue Variation by Fare Class Type and Cabin Accommodation

Concept	Expected	Simulations
Total Revenue (premium + economy)	Increases	Increases slightly/Decreases
FC 1-4 Revenue	Unchanged	Decreases
FC 5-10 Revenue	Increases	Increases
Premium Cabin Revenue	Increases	Increases
Economy Cabin Revenue	Unchanged	Decreases

Table 5-3: Comparison of Expected and Observed Results – Full EMSR

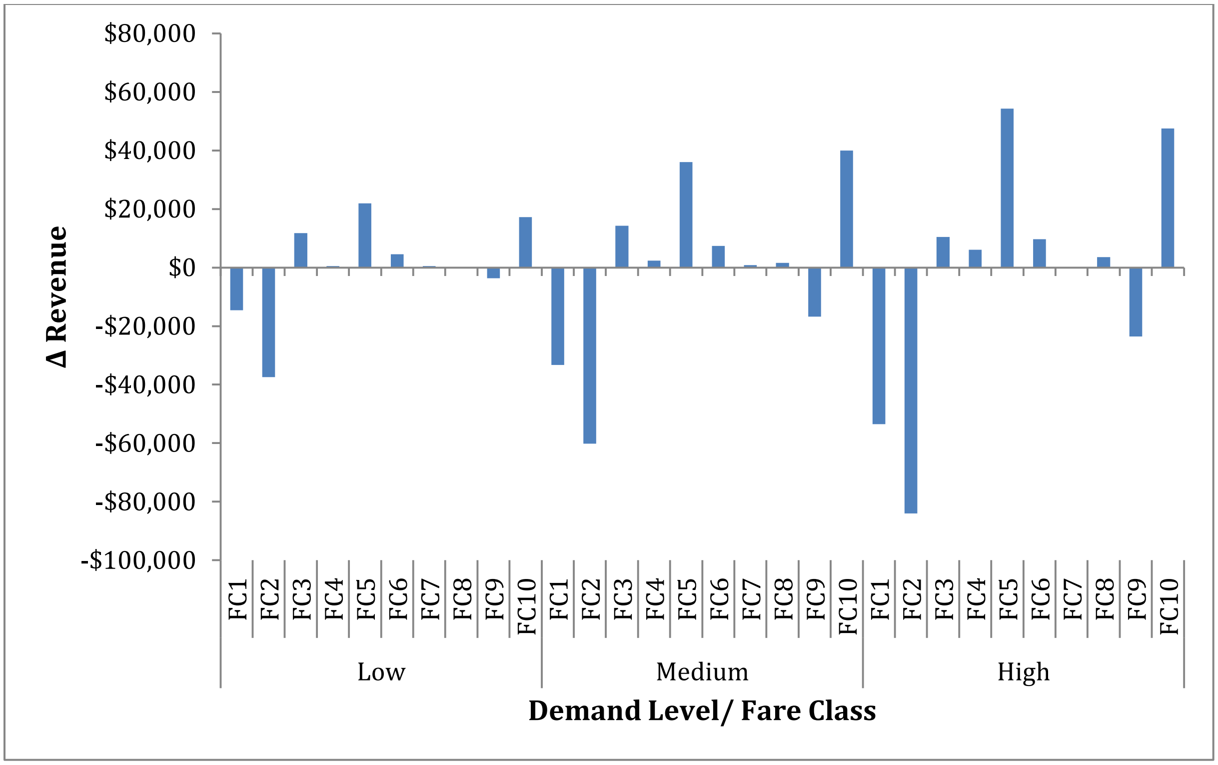

Figure 5-5: Full EMSR - Revenue Variation by Fare Class and Demand Level

In order to interpret and analyze these results, it is useful to revisit the principles of RM presented in the previous sections: "… RM is the process that determines the number of seats to be made available to each fare class on a flight…". Moreover, its main objective is to protect seats for later-booking, high-fare business passengers by forecasting the expected booking demand and mathematically optimizing to determine the number of seats to be protected from lower fare classes (Belobaba et al., 2015).

It should be noted that the absolute revenue increases or decreases observed for a fare class correspond to an increase or decrease, respectively, in bookings in such fare class. Applying the RM definition provided in the last paragraph to the presented results, it would be said that when Full EMSR is implemented:

- Fewer seats are protected for FC1 and FC2 (the highest fare classes available).
- More seats are protected for FC5 (the highest economy fare class available).
- More seats are left available for FC10 or, similarly, less seats are protected for all the other higher fare classes.

Based on the definitions provided above, a decrease in the protection levels for FC1 and FC2 is the result of an optimization that is fed with reduced forecasts for such fare classes. Forecasts are based on the previous observations of bookings by the airline; as described in Chapter 3, in PODS the booking information observed after each sample is added to the historical booking database of the RMS of the airline. In this case, since premium cabin seats are made available for economy fare classes, there is a potential for increasing the booking requests accepted in FC5 to FC10. In addition, since the space in premium cabin is limited, it is also possible that as a result of allowing bookings in economy fare classes to use seats in premium cabin, a booking request for FC1 or FC2 would not be able to be accepted 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 in future flights, reducing the number of seats to be protected, and so on, in a cyclical process.

Table 5-4 presents, on average across trials, the proportional change of the forecasts for FC1, FC2 and FC5, between the Full EMSR simulations and the baseline scenario for each of the simulated demand levels. This table shows two interesting results: on the one hand, that the forecasts for these fare classes are consistently lower for FC1 and FC2 when premium cabin capacity is shared with economy fare classes. On the other hand, the proportional decrease in the forecasts for these premium fare classes increase when the demand increases; therefore, the higher the demand, the larger the difference between the bookings forecasted with the distinct cabin and the premium cabin capacity sharing methodologies. In contrast, there is an increase in the FC5 forecasts as a result of the increase number of observed bookings in such fare class, as additional capacity is made available when premium cabin capacity is shared. In addition, it also shows that the proportional increase of the forecast of FC5 bookings when premium cabin capacity is shared compared to the distinct cabin scenario increases with the demand level.

Fare Class	Demand Level	Forecast Distinct EMSRb	Forecast Full EMSR	Variation
FC1	Low	123.49	102.27	-17.2%
	Medium	167.02	115.08	-31.1%
	High	218.62	126.97	-41.9%
FC2	Low	605.97	515.50	-14.9%
	Medium	747.82	563.29	-24.7%
	High	882.30	585.49	-33.6%
FC5	Low	1308.46	1341.24	2.5%
	Medium	1457.03	1524.65	4.6%
	High	1596.83	1701.68	6.6%

Table 5-4: Full EMSR - Proportional Variation of Forecasts – Leg-based

Being one of the main inputs for the optimizer, 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. So, considering the decrease in the forecasts for the highest premium fare classes, less protection from the lower premium fare classes is applied. In the case of Fare Class 1, for example, reduced levels of protection from lower fare classes results in a higher closure rate (Figure 5-6); 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. It should be noted as well that the gap between the FC1 closure rate of EMSRb and Full EMSR increases as the demand increases.

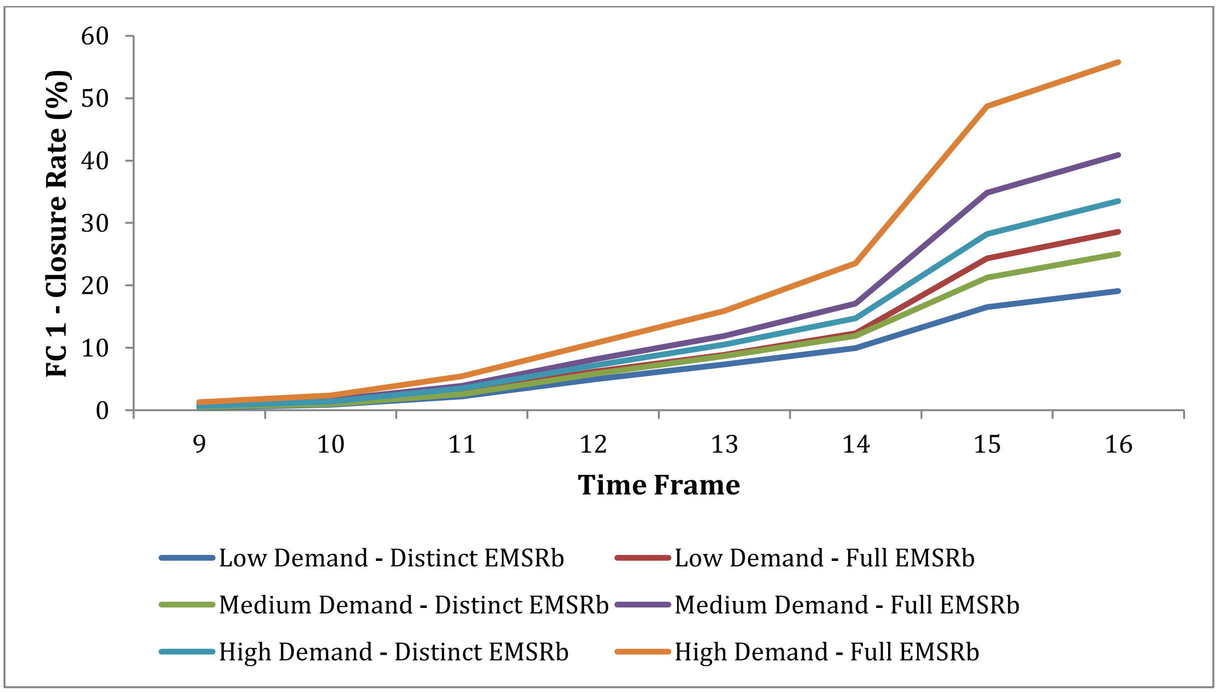

Figure 5-6: Closure Rate Fare Class 1 – Time Frames 11 to 16

The case of FC5 is different as additional seats are being made available for bookings in that class and the closure rates decrease accordingly (Figure 5-7). Therefore, considering the increase in the forecasts for FC5 more protection from the lower premium fare classes is applied, allowing increasing the bookings in the late time frames of the booking process. Finally, Figure 5-8 shows that the closure rates for FC10 are also lower when premium cabin capacity is shared, and explains why there is a significant increase of revenue in FC10.

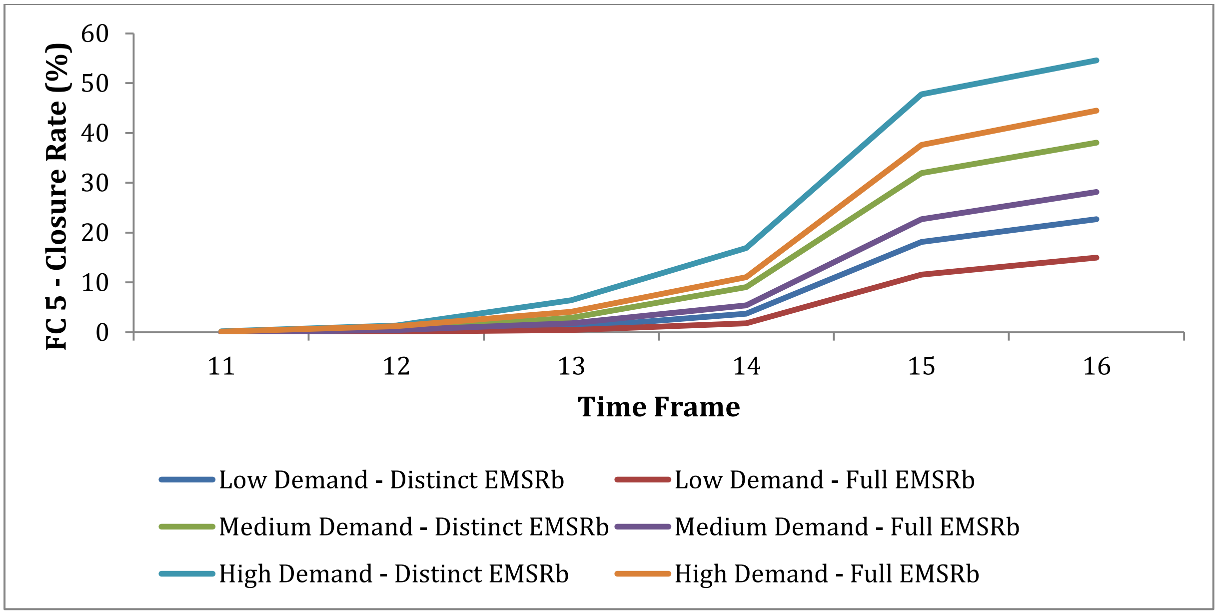

Figure 5-7: Closure Rate Fare Class 5 – Time Frames 11 to 16

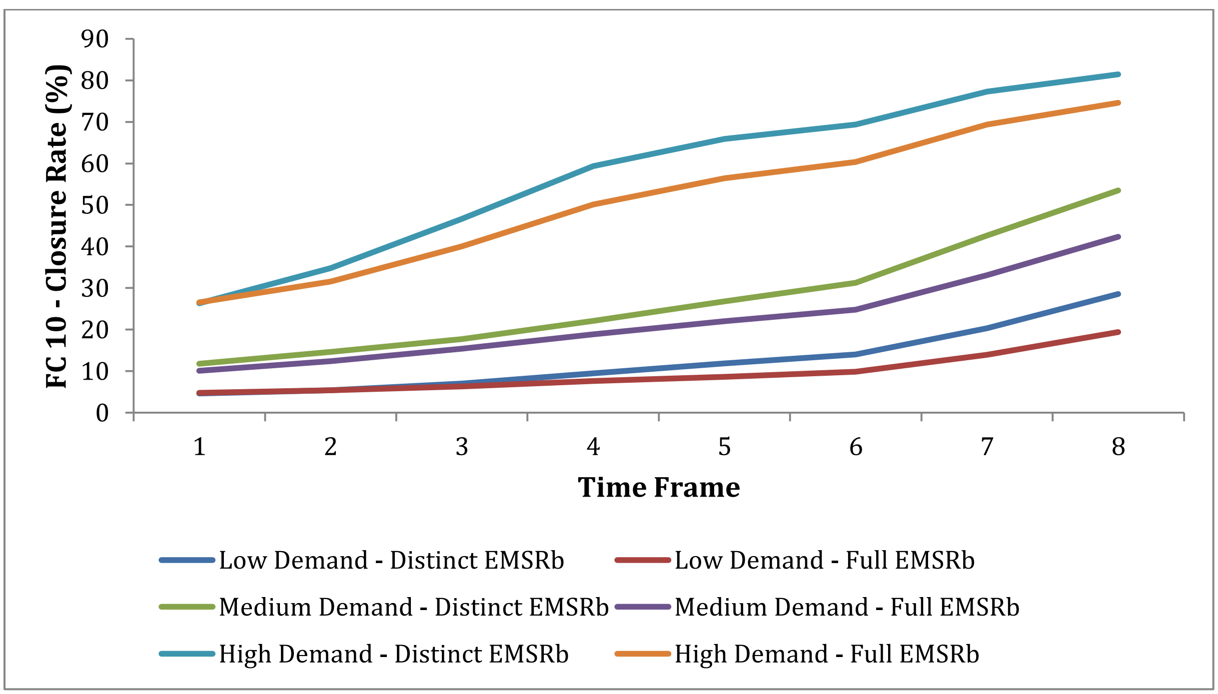

Figure 5-8: Closure Rate Fare Class 10 – Time Frames 1 to 8

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

As discussed in the initial sections of this chapter, the motivation for implementing a premium cabin capacity sharing mechanism is to increase the revenue captured by accommodating additional passengers that buy economy fare classes in premium cabin seats that otherwise would have been flown empty. Conceptually, these additional passengers should not replace or displace the passengers wishing to fly in premium cabin; however, the results of the simulations have shown that Full EMSR results in a reduction in the revenue captured from FC's 1 to 4. Despite the increase in total revenue (premium + economy), airlines value greatly the passengers that buy high value premium fare classes and a reduction in that source of revenue is something that the management teams of airlines that operate dual cabin aircraft usually find problematic and not strategic in the long term.

Taking this into consideration, two mechanisms for protecting premium fare classes' revenue are proposed. The first of these consists in delaying premium cabin capacity sharing until a predetermined time frame of the booking process with the intention of guaranteeing that capacity is only shared only with passengers that buy some of the high-value economy fare classes with little or no Advance Purchase (AP) requirements. The second scheme proposed consists of activating premium cabin capacity sharing on a leg only after a certain predefined economy cabin load factor has been reached; this approach is also useful to avoid sharing premium capacity throughout the booking process and guarantees that premium cabin capacity sharing is restricted only to flights with a considerable economy cabin load factor.

- **Time Frame Protection**

Time Frame Protection ("TFP") is tested using three different predetermined time frames:

- Time Frame 10 (TFP=10): distinct until 14 days before
- Time Frame 12 (TFP=12): distinct until 7 days before departure
- Time Frame 14 (TFP=14): distinct 3 days before departure

For all these cases, it is found that the highest total revenue gains are found when premium cabin capacity is made available for passengers with bookings in economy fare classes just three days before departure. While in the low demand scenario the gains increase from 0.03% to 0.51%

(Figure 5-9), in the medium and high demand scenarios TFP turns losses into gains. As a matter of fact, although the proportional losses obtained in the high demand scenario without TFP are larger than in the medium demand scenario, the highest proportional revenue gains are achieved in the high demand scenario with TFP=14 (1.06%) (Figure 5-10 and Figure 5-11). 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; an example of this is provided in Figure 5-12 for the medium demand scenario.

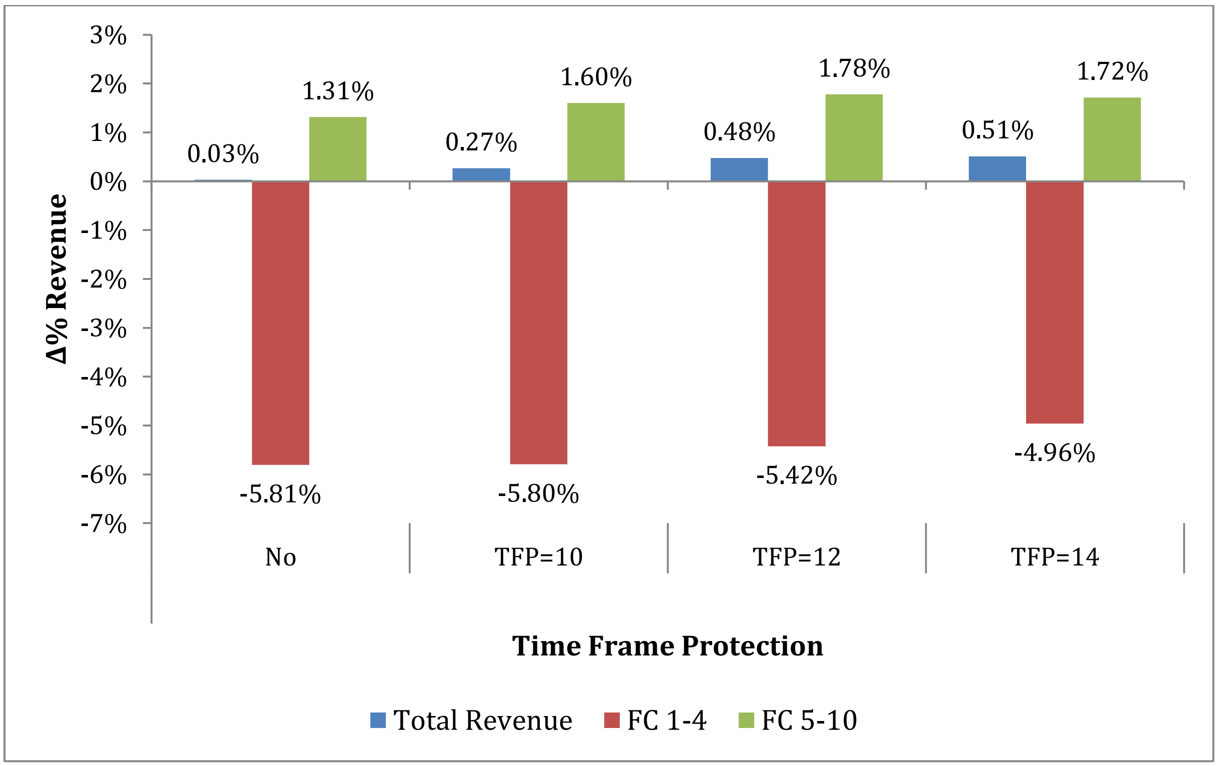

Figure 5-9: Time Frame Protection – Full EMSR – Low Demand

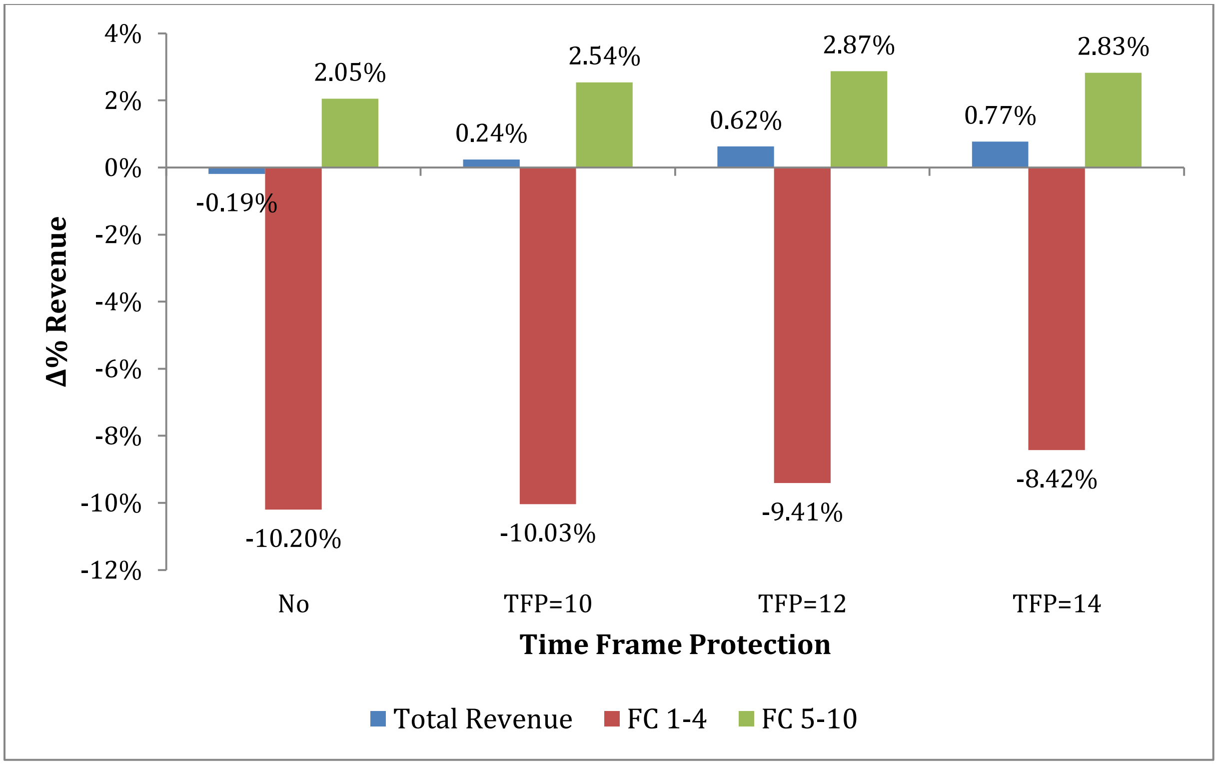

Figure 5-10: Time Frame Protection - Full EMSR - Medium Demand

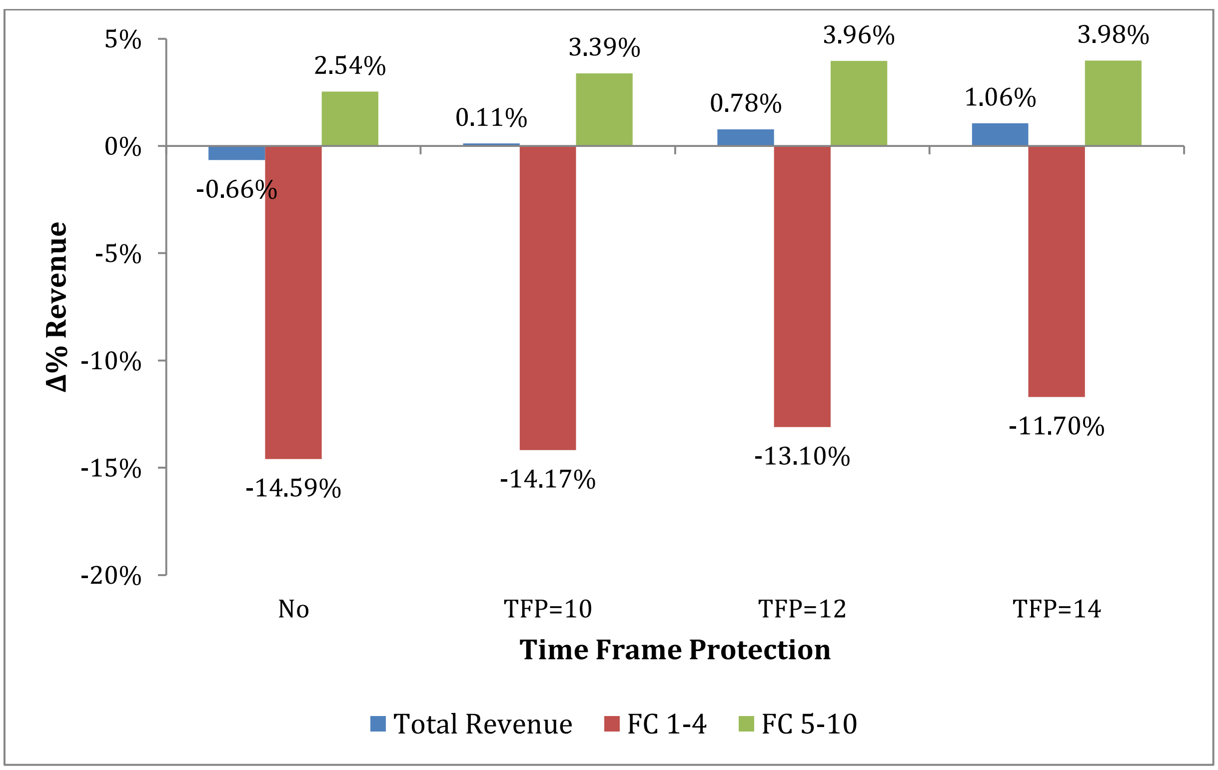

Figure 5-11: Time Frame Protection - Full EMSR - High Demand

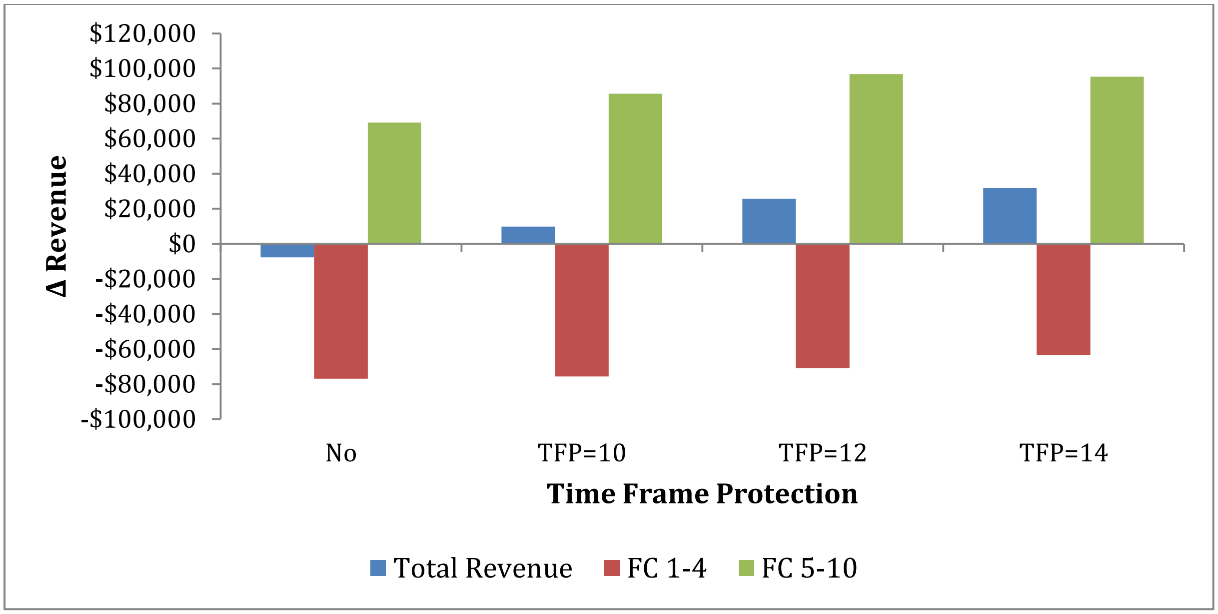

Figure 5-12: Full EMSR with TFP - Absolute Revenue Gains by Fare Class Type – Medium Demand

Two other perspectives are offered to understand the effect of TFP in AL1's revenues when using Full EMSR; the first of these consists in analyzing how the bookings are allocated between the cabins. Figure 5-13 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, this is just part of the explanation for the changes. The main source of revenue gains when TFP=14 comes from passengers buying economy fare classes and effectively flying in economy cabin; in the referenced figure, for example, this category passes from losing $42,275 to gaining $34,192 in the medium demand scenario.

The second perspective to understand the impact of time frame protection is given by the variations in revenue by fare class, as presented in Figure 5-14: it is evident that the largest variations are observed in FC5 and FC10. By simply activating TFP, FC10 revenue decreases considerably (even since TFP=10) while all the other economy fare classes increase the revenue; therefore, the conclusion is that sharing availability since the beginning of the booking process allows for a large number of bookings in FC10 and that when availability is not shared, many of those passengers need to book their tickets in higher fare classes.

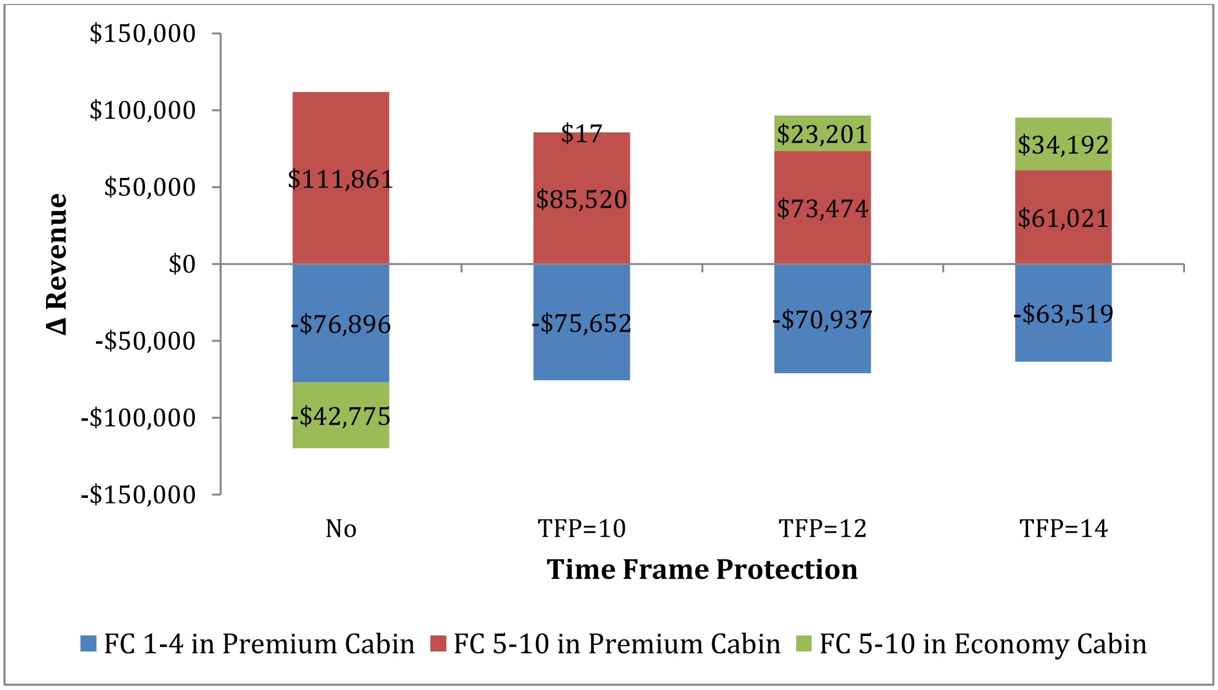

Figure 5-13: Full EMSR - Absolute Revenue Variation by FC Type, Cabin Accommodation and TFP – Medium Demand

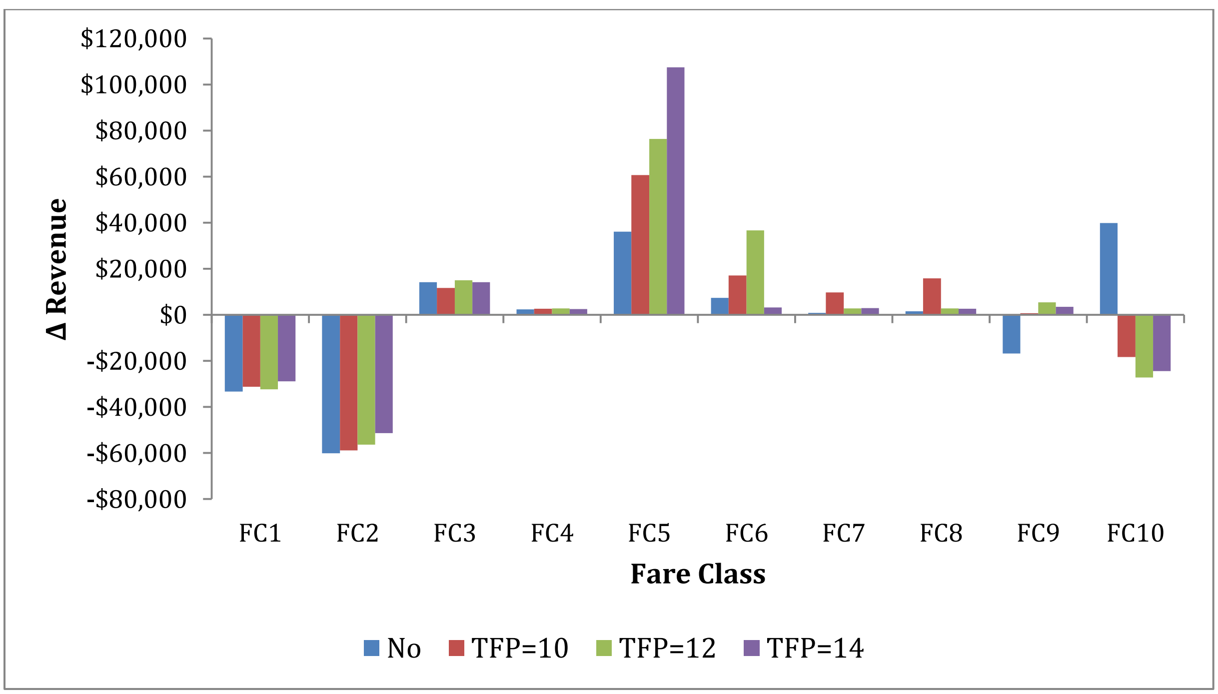

Figure 5-14: Absolute revenue variations by fare class – Full EMSR with TFP – Medium Demand

- **Leg Load Factor Criterion Protection**

The Leg Load Factor Criterion ("LLFC") is another mechanism that is used to delay the activation of premium cabin capacity sharing in a leg until its associated load factor exceeds a predetermined leg load factor. LLFC is tested with Full EMSR at the following levels: 70%, 80%, 90%, 95% and 99%.

Figure 5-15, Figure 5-16 and Figure 5-17 present the results for Full EMSR with different values of LLFC simulated at the low, medium and high demand levels, respectively. Similar to what was found using Time Frame Protections, implementing LLFC turns losses (or negligible revenue gains) into considerable proportional total revenue gains; in addition, the proportional losses in premium fare classes are reduced, as it was the intention with the implementation of the protection mechanisms. Finally, it is found that a 95% LLFC levels helps to achieve the highest total revenue gains.

The distribution of the gains and losses between fare class groups and cabins changes in a similar way as it does when TFP is used: with high LLFC values there is less premium cabin capacity shared, so the revenue gains from passengers that book in FC 5 to 10 but are carried in premium cabin decrease together with the losses from passengers that book in FC 1 to 4. However, at the levels where the highest revenue gains are reported, there is a significant increase in the revenue captured from passengers flying in economy cabin with respect to the baseline (Figure 5-18).

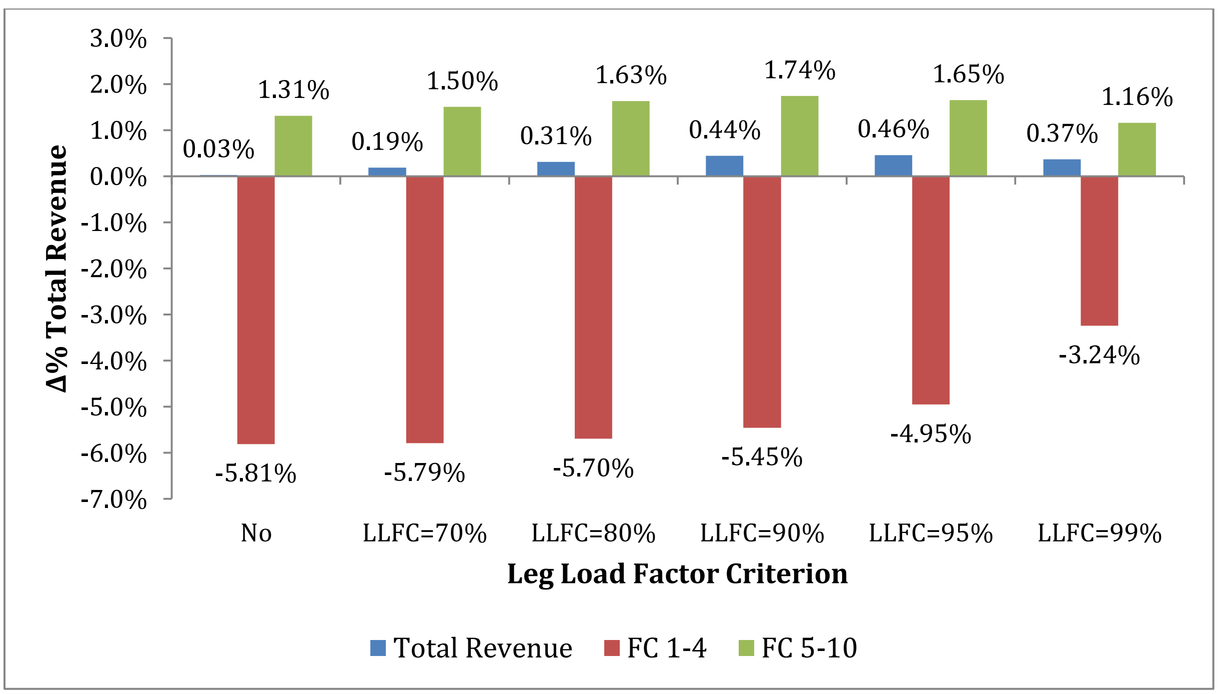

Figure 5-15: Leg Load Factor Criterion – Full EMSR - Low Demand

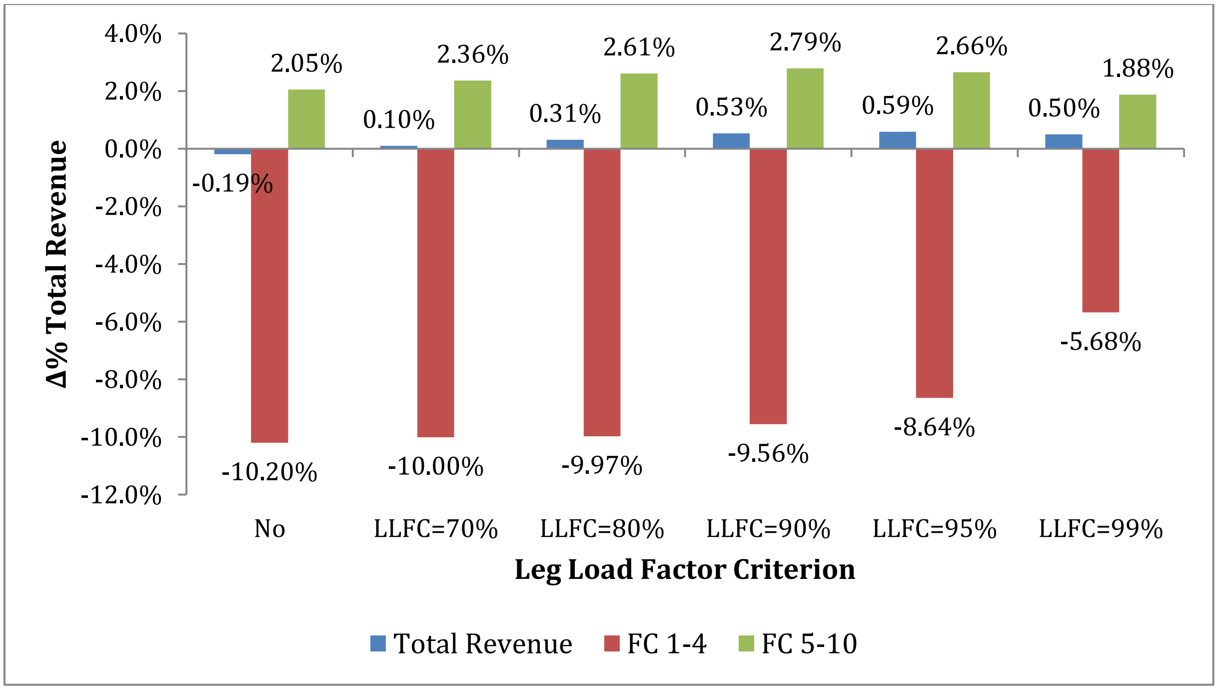

Figure 5-16: Leg Load Factor Criterion – Full EMSR - Medium Demand

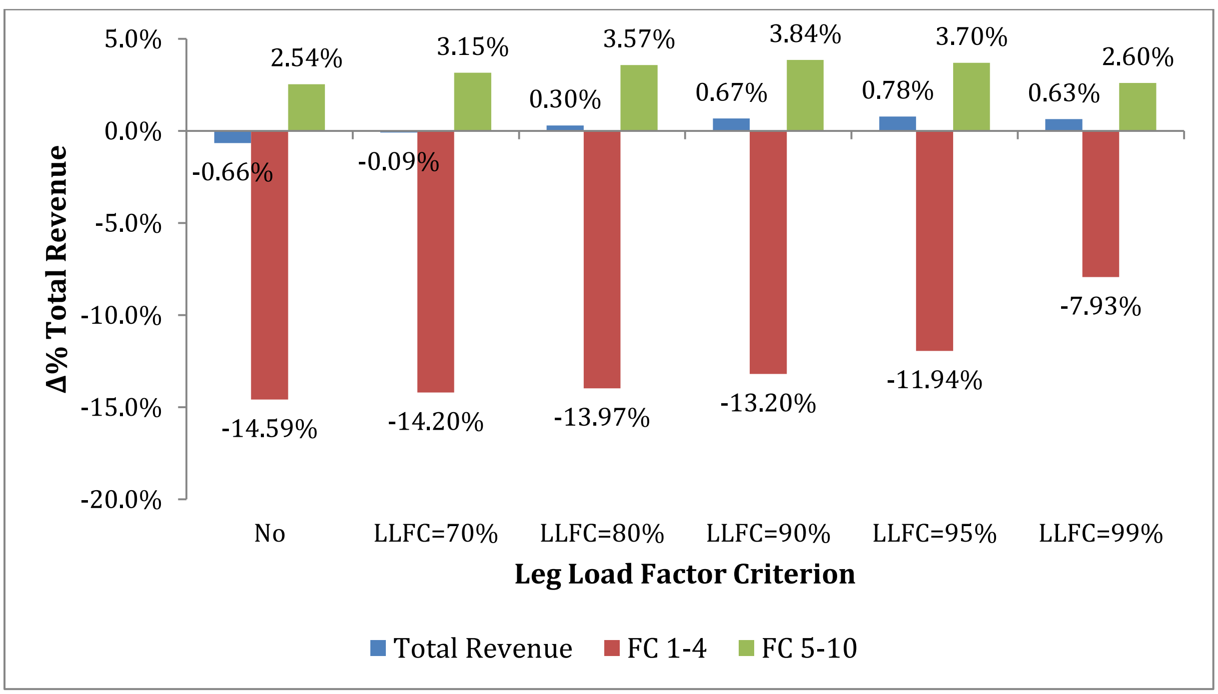

Figure 5-17: Leg Load Factor Criterion – Full EMSR - High Demand

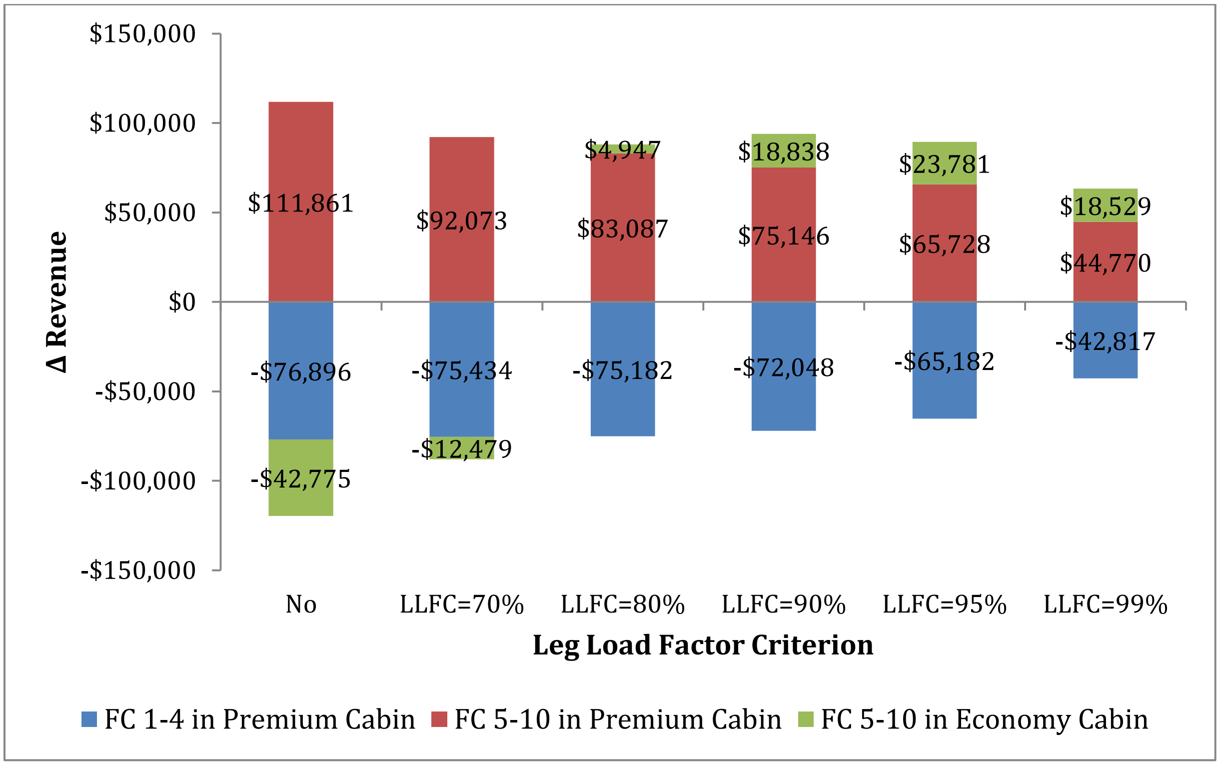

Figure 5-18: Full EMSR - Absolute Revenue Variation by FC Type and Cabin Accommodation and LLFC – Medium Demand

5.2 OD-Control RM Premium Cabin Shared Capacity Schemes

As described in Chapter 3, three heuristics are proposed as mechanisms to increase the total revenue of an airline that operates dual cabin aircraft and uses OD RM Control methods; each of the following sections explores the results for each of these heuristics.

The scenario used in PODS as a baseline to compare the performance of each of the capacity sharing methods is presented in Table 5-5. Running on Network V1, AL1, AL2 and AL4 use OD RM, specifically with ProBP as the optimizer. Furthermore, they manage the capacity of each of the two cabins (premium and economy) independently, as if each cabin was a different aircraft. Meanwhile, AL3 uses a Leg-Based RM method, specifically EMSRb.

OD-Control Base Case
Network V1
AL1: Distinct ProBP - Path-based Std. Forecasting
AL2: Distinct ProBP – Path-based Std. Forecasting
AL3: EMSRb – Leg- Based Std. Forecasting
AL4: Distinct ProBP – Path-based Std. Forecasting

Table 5-5: OD RM Methods Baseline

5.2.1 Method 1: Premium Cabin Bid Price Capacity Sharing

The simulation settings for evaluating the impact of the implementation of Premium Cabin Bid Price Capacity Sharing (hereinafter referred to as "Method 1" or "PCBP") are described in Table 5-6. The modification with respect to the OD RM Methods Baseline is that Airline 1 shares its premium cabin capacity with economy fare classes based on the rules of the first mechanism described in Section 3.5.2; briefly, premium cabin capacity is made available to economy fare classes when the economy cabin is sold out and the fare of an economy fare class is higher than the bid price applicable for the premium cabin (or for connecting markets, higher than the sum of the premium cabin bid prices of the legs traversed by the ODF).

OD-Control Premium Cabin Bid Price Capacity Sharing
Network V1
AL1: Shared ProBP
Premium Cabin Bid Price Capacity Sharing (Method 1)
Path-based Std. Forecasting
AL2: Distinct ProBP – Path-based Std. Forecasting
AL3: Distinct EMSRb – Leg- Based Std. Forecasting
AL4: Distinct ProBP – Path-based Std. Forecasting

Table 5-6: Premium Cabin Bid Price Capacity Sharing - Simulation settings

Figure 5-19 presents the impact of the implementation of Method 1 by AL1 on the total revenue (premium plus economy) of each of the airlines competing in the network at different demand levels. AL1 proportional gains range from 0.54% to 1.08%, depending on the demand level (the higher the demand, the higher the proportional gain). In contrast, all the competitors see their total revenue reduce as a consequence of the decision of AL1 to share premium cabin capacity with economy fare classes on its own flights (i.e. fare classes 5 to 10).

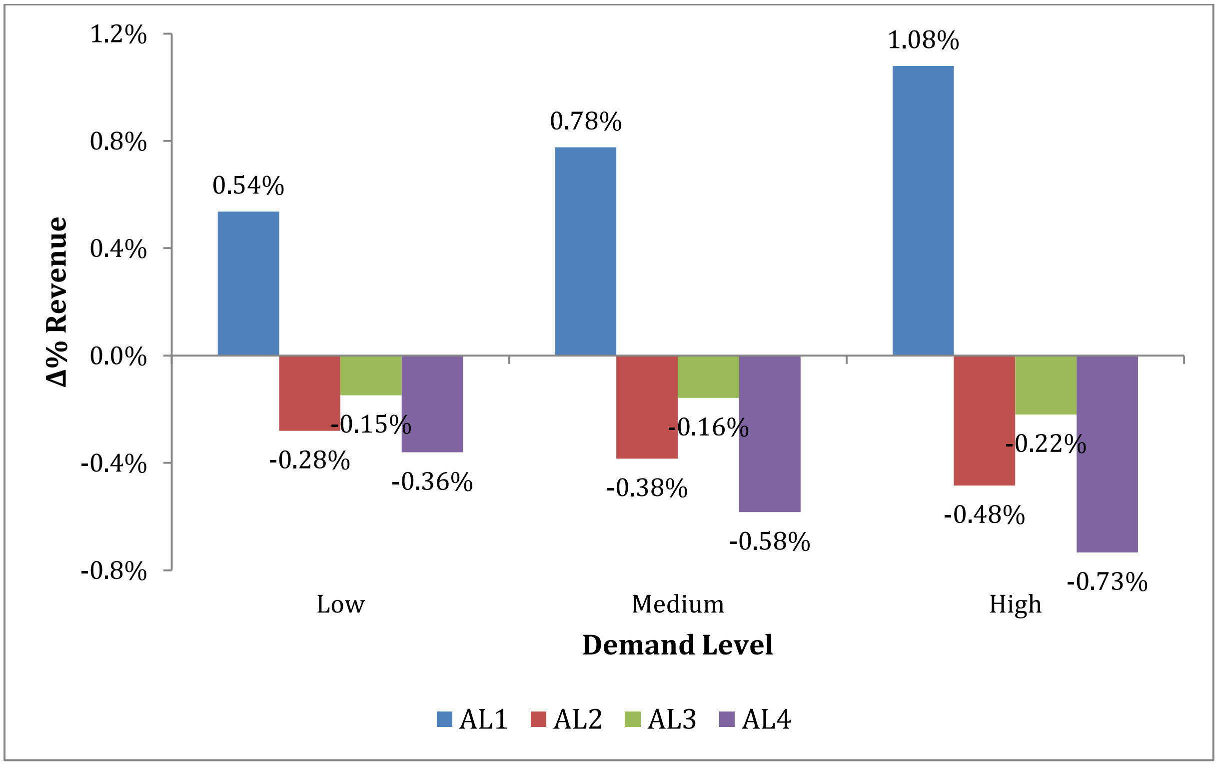

Figure 5-19: Method 1 - Total Revenue Proportional Variation by Airline

Observing an increase in AL1's revenue and a decrease in the revenue of the other airlines is an expected result of an implementation of a shared capacity strategy such as the one proposed herein. Intuitively, if the idea is simply to share the remaining capacity in premium cabin with additional passengers of economy fare classes, what would be expected is an increase in total revenue resulting from the growth of the revenue captured from economy fare classes (due to the passengers accommodated in premium cabin) while keeping the same revenue from premium fare classes.

Overall, it is true that this mechanism effectively generates such total revenue gains; it is also true that the revenue from economy fare classes increases due to the additional space available. However, the expectation of premium fare classes not being affected by the implementation of shared capacity is false. As shown in Figure 5-20, the total revenue gains resulting from sharing capacity are a combination of gains (1.60% to 3.63%) from economy fare classes with losses from premium fare classes (-4.15% to -9.91%). It should be noted that despite the magnitude of such proportional losses in premium fare classes, an increase in total revenue is achieved at all the demand levels because premium fare classes' revenue represents less than 19% of the total revenue of AL1. Therefore, a reasonably low proportional increase in such economy fare classes compensates for proportional losses of larger magnitude in premium fare classes; this is clarified in Figure 5-21, which presents the absolute gains and losses for each of the fare class types and explains the total revenue gains.

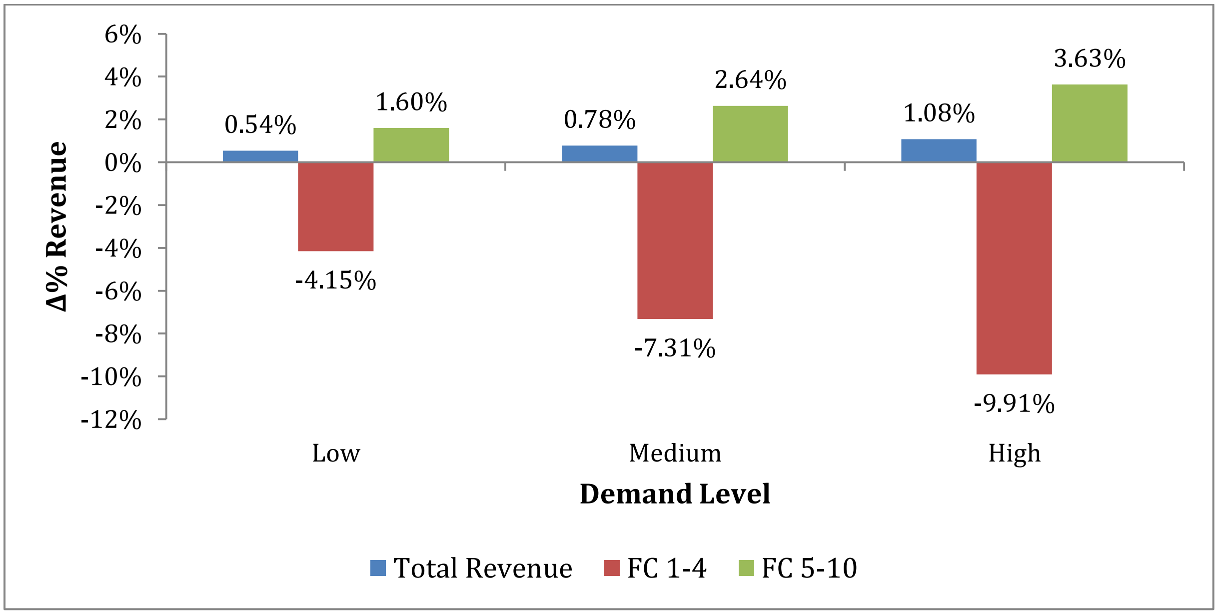

Figure 5-20: Method 1 - Proportional Revenue Gains by Fare Class Type

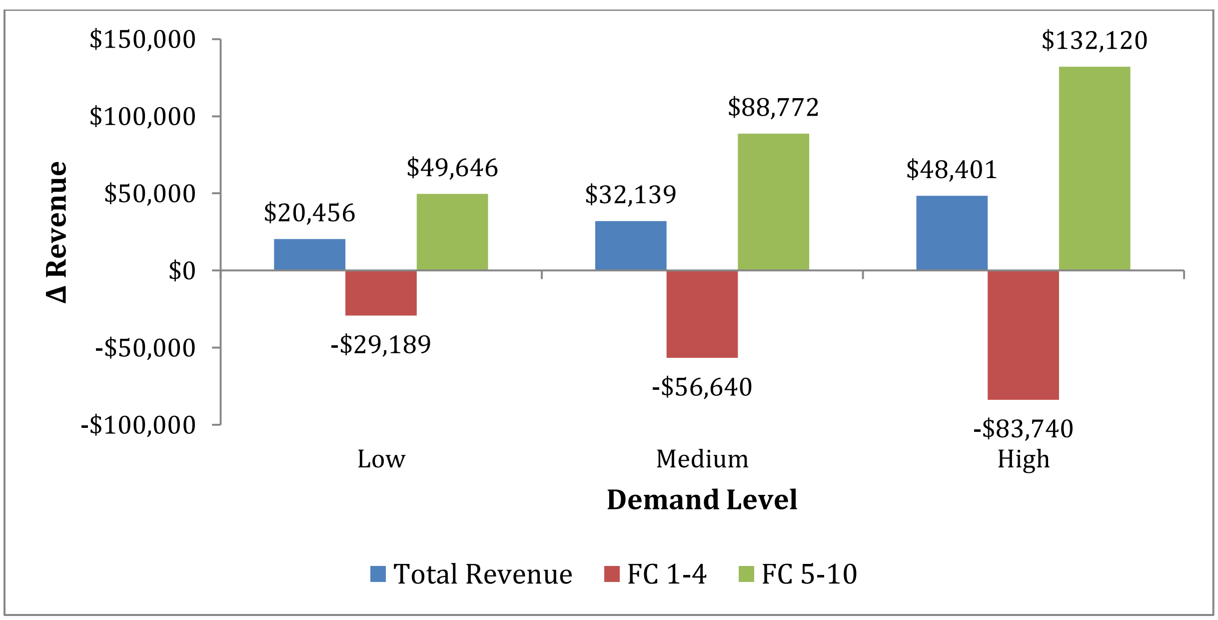

Figure 5-21: Method 1- Absolute Revenue Gains by Fare Class Type

Relying again on intuition, the revenue gains in FC 5-10 would be expected only as a result of the additional passengers accommodated buying such fare classes that would be able to be accommodated in the empty seats of premium cabin. As a matter of fact, considering that the Premium Cabin Bid Price methodology for capacity sharing only takes effect once economy cabin is sold out, it would be expected that the revenue obtained from the economy cabin should not be affected at all by the implementation of premium cabin capacity sharing; once more, the results are different from what would have been expected.

Figure 5-22 reports the absolute revenue variations by fare class type (i.e. premium or economy fare classes) and the cabin where the passengers were finally accommodated. It is found that the increase in economy fare classes (FC 5 to 10) is split between gains from passengers flying in premium cabin (as expected) and gains from passengers flying in economy cabin (not expected intuitively). Considering that implementing capacity sharing has the intention of making the premium cabin more productive by accommodating additional passengers in it, Figure 5-22 presents two surprising results: first, that implementing premium cabin capacity sharing can make the economy cabin more productive by increasing the revenue captured from it. Second, that capacity sharing does not necessarily generate more revenue from its premium cabin: as observed, while the revenue captured from passengers flying in premium cabin increases by $4,977 in the low demand scenario ($34,166 of increase from FC5-10 flying in premium cabin minus $29,189 of decreases in FC1-4 revenue), it decreases at the medium and high demand scenarios by $2,687 and $15,271, respectively. All these differences between the expected and observed results are summarized in Table 5-7.

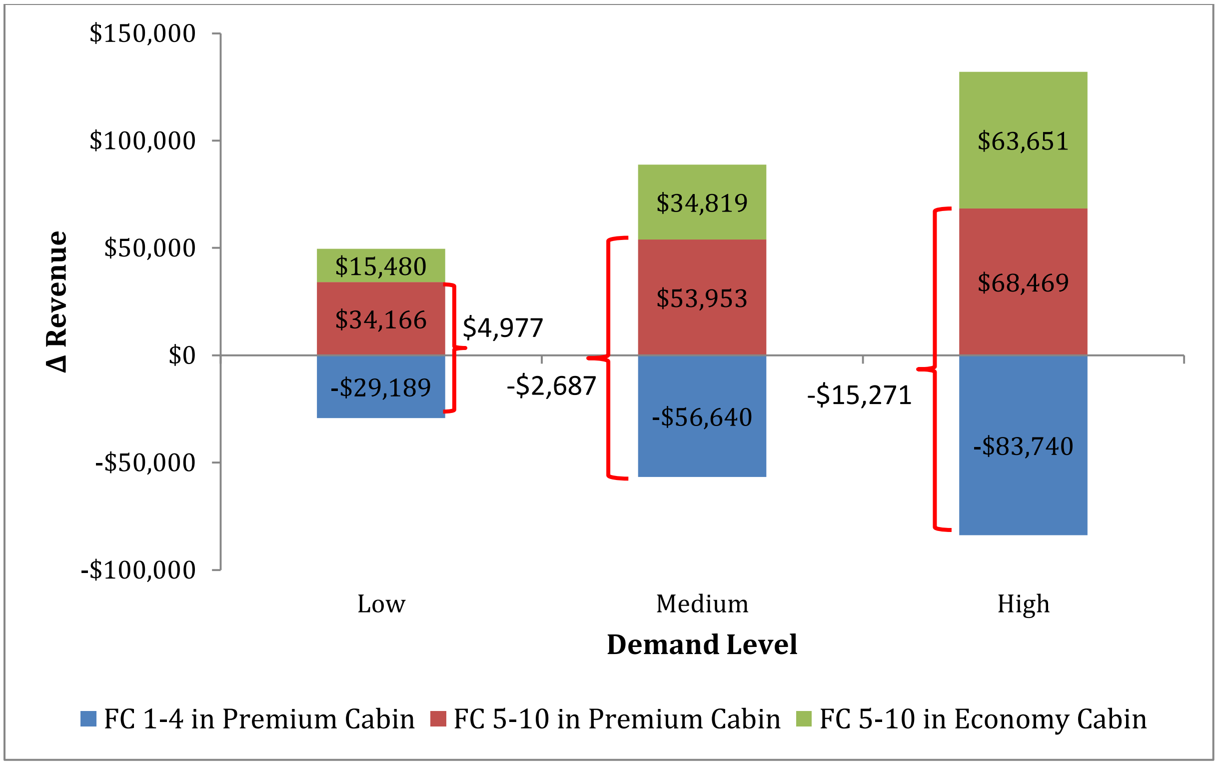

Figure 5-22: Method 1 - Absolute Revenue Variation by Fare Class Type and Cabin Accommodation

Concept	Expected	Simulations
Total Revenue (premium + economy)	Increases	Increases
FC 1-4 Revenue	Unchanged	Decreases
FC 5-10 Revenue	Increases	Increases
Premium Cabin Revenue	Increases	Mixed – Sometimes increases, sometimes decreases
Economy Cabin Revenue	Unchanged	Increases

Table 5-7: Comparison of Expected and Observed Results - PCBP Capacity Sharing Method

In order to understand the changes in the revenue obtained from each of these groups, we can analyze the results from a more granular perspective, as provided by the changes in absolute revenue of each fare class at each demand level, presented in. As shown in Figure 5-23, the revenue variations are mainly explained by revenue reductions in FC1, FC2 (the two highest premium fare classes) and FC10 (the lowest fare class) compensated by revenue gains in FC5 (the highest among the economy fare classes); although all the other fare classes vary as well, the variations are negligible when compared to the changes in the revenue of the fare classes mentioned. It is also observable that the higher the demand level, the larger the magnitude of the respective gains or losses.

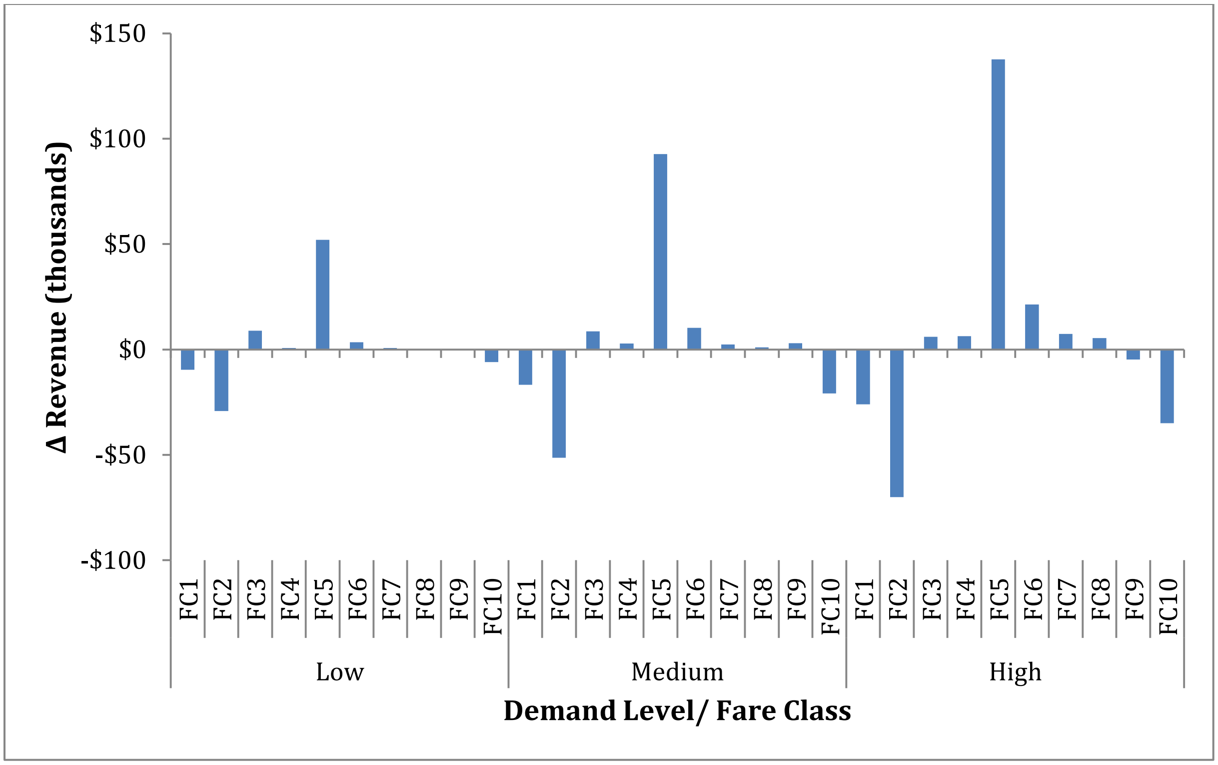

Figure 5-23: Method 1 - Revenue Variation by Fare Class and Demand Level

As explained in the Full EMSR case, a decrease in the protection levels for FC1 and FC2 is the result of an optimization that is fed with reduced forecasts for such fare classes. Table 5-8 presents, in average across trials, the proportional change of the forecasts for FC1, FC2 and FC5, between the Method 1 simulations and the baseline scenario for each of the simulated demand levels. As with the Full EMSR case, it is found that the forecasts for fare classes 1 and 2 are consistently lower when premium cabin capacity is shared with economy fare classes. Again, the proportional reduction in the forecasts for these premium fare classes is larger when the demand increases. In the case of FC5, its forecast increases as a result of the increase number of observed bookings in such fare class, as additional capacity is made available when premium cabin capacity is shared. In addition, it also shows that the proportional increase of the forecast of FC5 bookings when premium cabin capacity is shared compared to the distinct cabin scenario increases with the demand level.

Fare Class	Demand Level	Forecast Baseline	Forecast Method 1	Variation
FC1	Low	92.8	85.1	-8.3%
	Medium	123.7	105.9	-14.4%
	High	170.4	133.2	-21.8%
FC2	Low	487.0	402.2	-17.4%
	Medium	645.4	494.5	-23.4%
	High	891.8	613.0	-31.3%
FC5	Low	831.8	878.5	5.6%
	Medium	940.0	1027.7	9.3%
	High	1061.2	1213.1	14.3%

Table 5-8: Method 1 - Proportional Variation of Forecasts

In the case of Bid Price OD Control, those variations are reflected in changes to the bid prices. So, considering the decrease in the forecasts for the highest premium fare classes, less protection from the lower premium fare classes is applied and therefore the premium cabin bid price tends to decrease when capacity is shared. Figure 5-24 compares the premium cabin bid prices with distinct control and with Method 1 for capacity sharing at the low, medium and high demand levels. It can be observed how the bid prices increase as demand increases, illustrating the effect of forecasted demand on the calculation of such bid prices; however, for the purposes of this thesis it is even more relevant to observe how the gap between the premium cabin bid prices compared in each figure also increases as the demand increases. Therefore, because of the impact of shared capacity in the forecasts for the highest fare classes, the decrease in protection from the lower premium fare classes (FC3 and FC4) is lower.

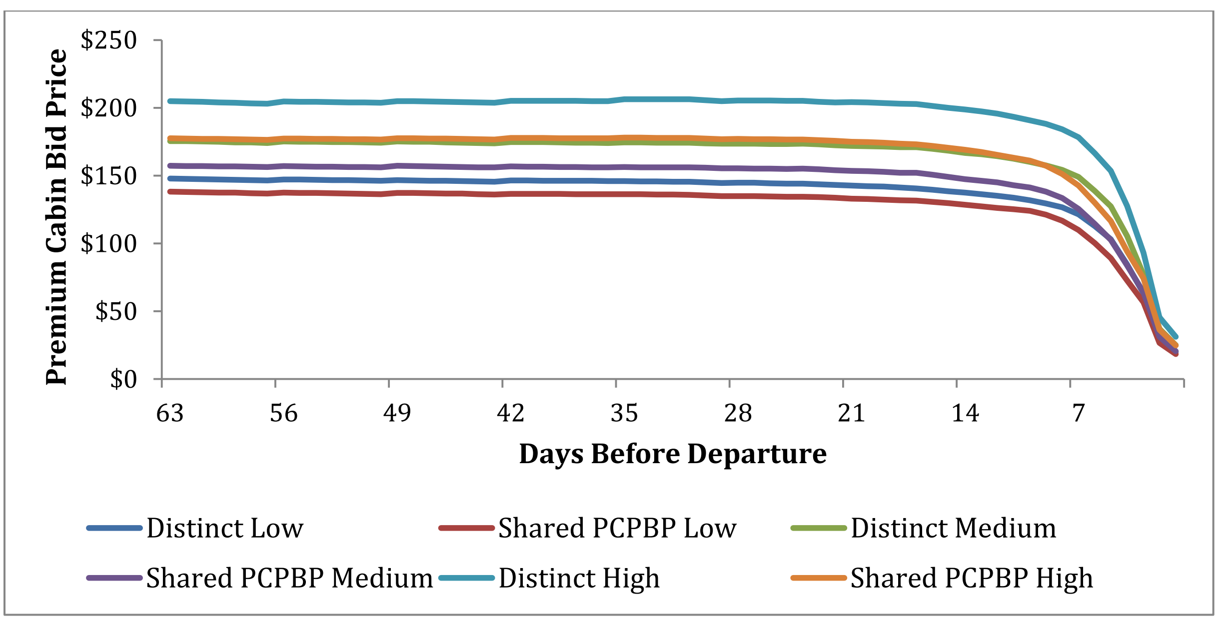

Figure 5-24: Premium Cabin Bid Prices - Distinct and Method 1

In contrast to the observed patterns of the premium cabin bid prices, the economy cabin bid prices tend to increase as a result of the higher forecasts observed for FC5, even though many of the additional FC5 bookings recorded are not using economy but premium cabin for travelling. Figure 5-25 illustrates the increase in the economy cabin bid price for each of the levels of demand simulated, and how the gap between the distinct and shared economy cabin BP's also grows as the demand grows, increasing protecting the highest economy fare classes, such as FC5 and FC6, from the lowest fare classes.

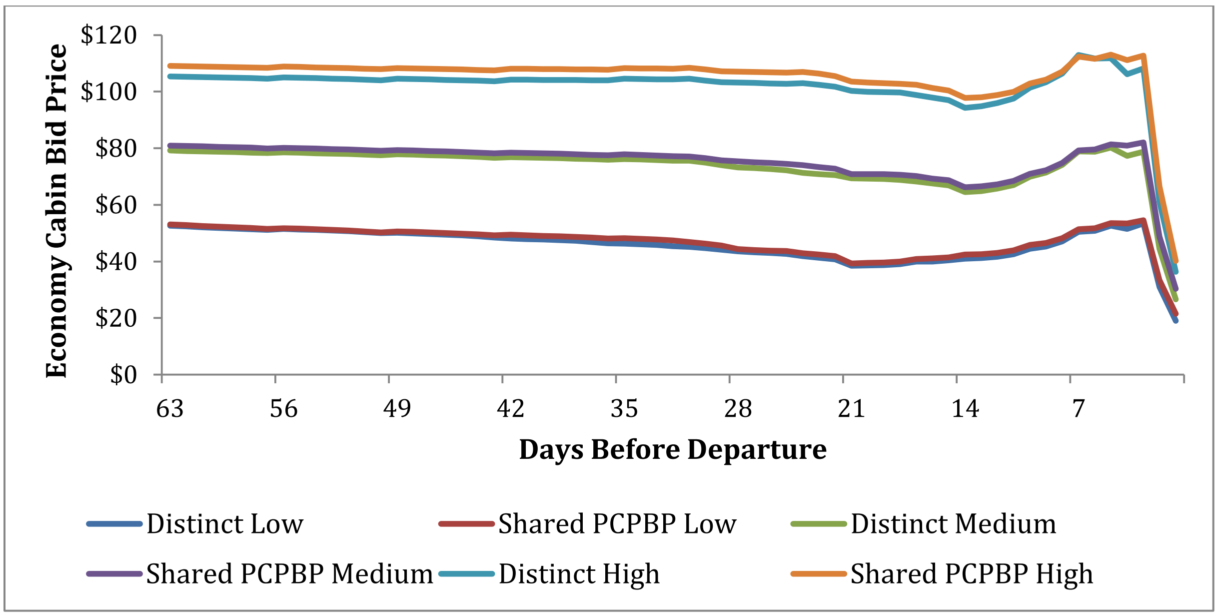

Figure 5-25: Economy Cabin Bid Prices - Distinct and Method 1

These changes in the bid prices affect how each fare class is opened or closed for each ODF; closure rates, measuring the percentage of ODF's in which each fare class is closed, show how each of such fare classes loses availability over the booking period. Based on the results of these simulations, the closure rates allow us to identify four groups of fare classes that share similar behaviors with respect to their closure rates within such groups, but also have different behaviors across groups: high premium fare classes (FC1 and FC2), low premium fare classes (FC3 and FC4), high economy fare classes (FC5 and FC6) and low economy fare classes (FC7, FC8, FC9 and FC10).

Analyzing since time frame 11 and until the end of the booking period, the FC1 closure rates are found to be higher when premium cabin capacity is shared than when each cabin is managed distinctly, at all demand levels. These closure rates are presented in Figure 5-26 and illustrate how, for example, the closure rate of FC1 increases from 23.1% to 27.6% in time frame 16 at the medium demand level, when premium cabin capacity is shared. This means that less booking requests in FC1 can be accepted in the last time frames, when actually most of such booking requests are expected; FC2 follows a similar pattern.

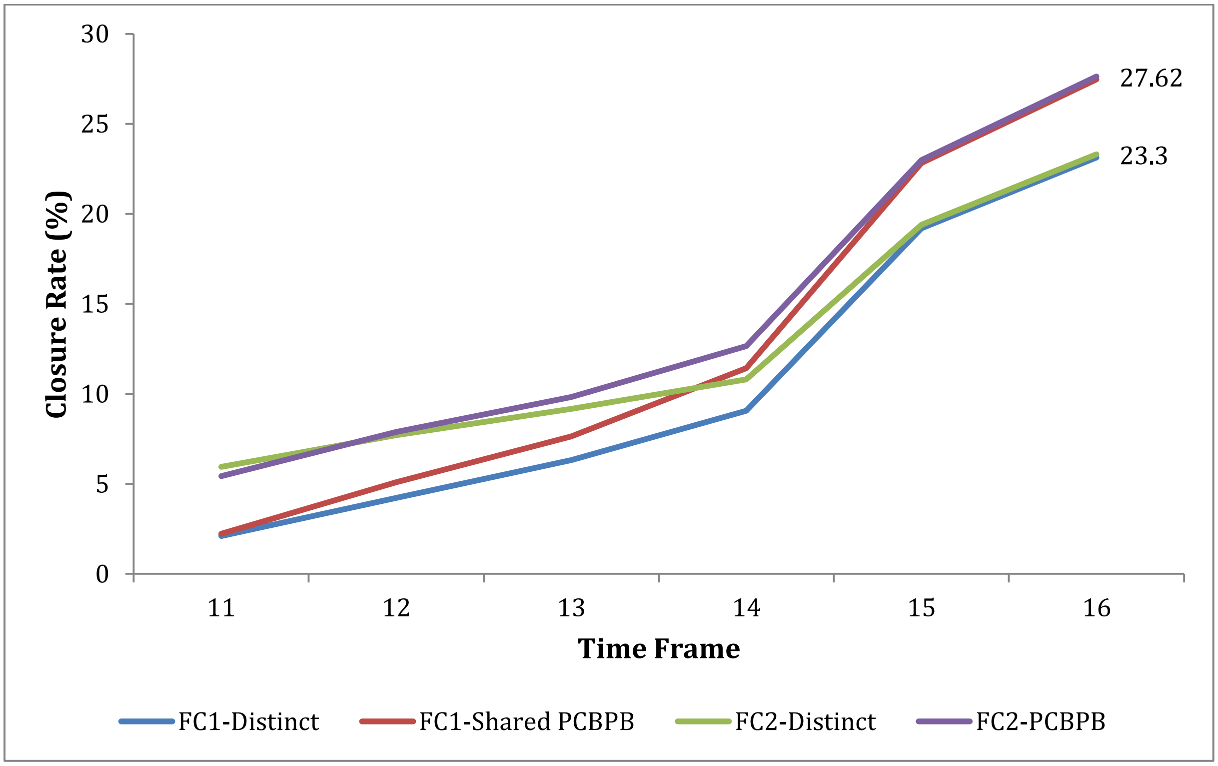

Figure 5-26: Closure Rates Fare Classes 1 and 2 - Time Frames 11 to 16 – Medium Demand

Higher closure rates when premium cabin capacity is shared for FC1 and FC2 have a significant effect on the number of bookings recorded by AL1 in such fare classes. As presented before, the losses in premium fare classes' revenue are fully explained by losses in FC1 and FC2; as illustrated in Figure 5-27 for the medium demand scenarios, such decrease in bookings for the high premium fare classes group is observed only since time frame 14 as a result of much higher closure rates.

The second group of fare classes examined is the group of the low premium fare classes, FC3 and FC4. The respective closure rates for the medium demand scenario are presented in Figure 5-28 and show that, when premium cabin capacity is shared, more flights keep accepting booking requests for such fare classes while these are still available (advance purchase "AP" requirements close 100% of them at time frame 9 and 13, respectively); these lower closure rates are the result of the lower premium cabin bid prices referenced before and allow to increase the number of booking requests that can be accepted for such fare classes, as presented in Figure 5-29 for the medium demand level case. While the impact of these additional bookings is relatively low considering the additional revenue captured specifically from such bookings, it is part of the explanation for the observed losses in FC1 and FC2; each seat booked in FC3 or FC4 is a seat that cannot be booked in the high premium fare classes, affecting historical bookings, forecasts and bid prices.

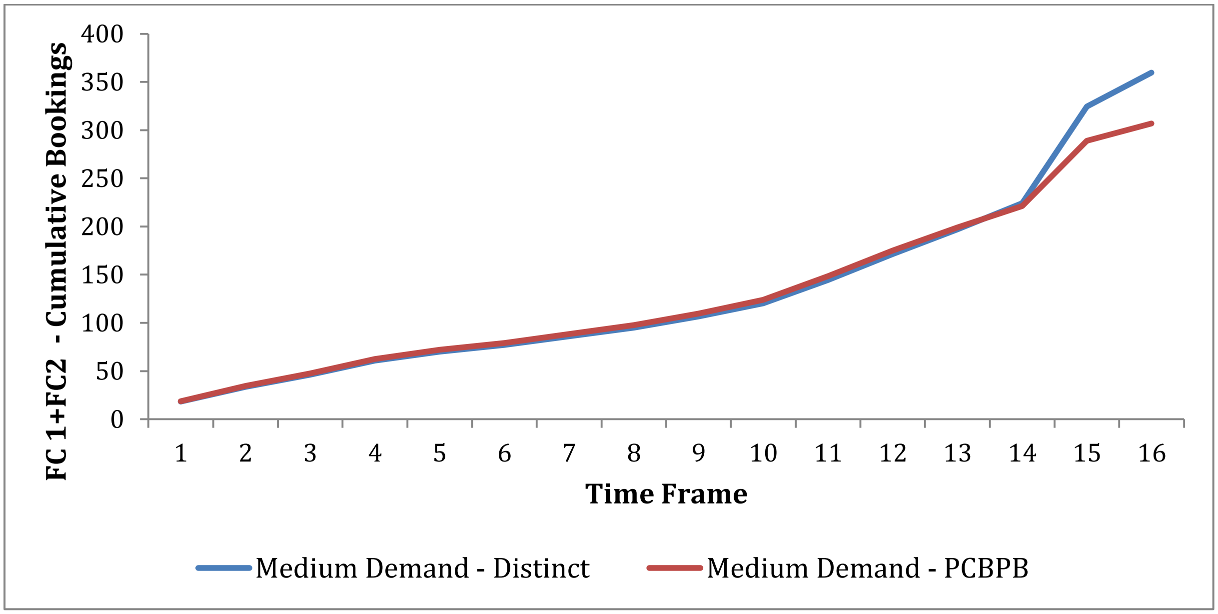

Figure 5-27: Cumulative Bookings FC1 + FC2 Medium Demand Level

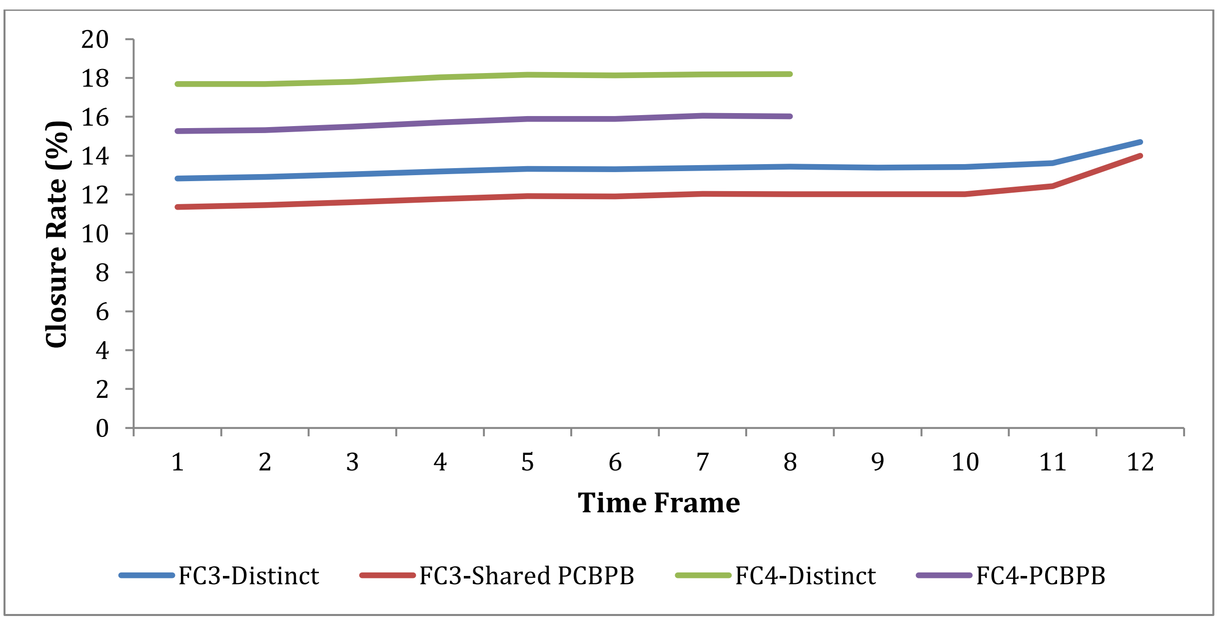

Figure 5-28: Closure Rates Fare Classes 3 and 4

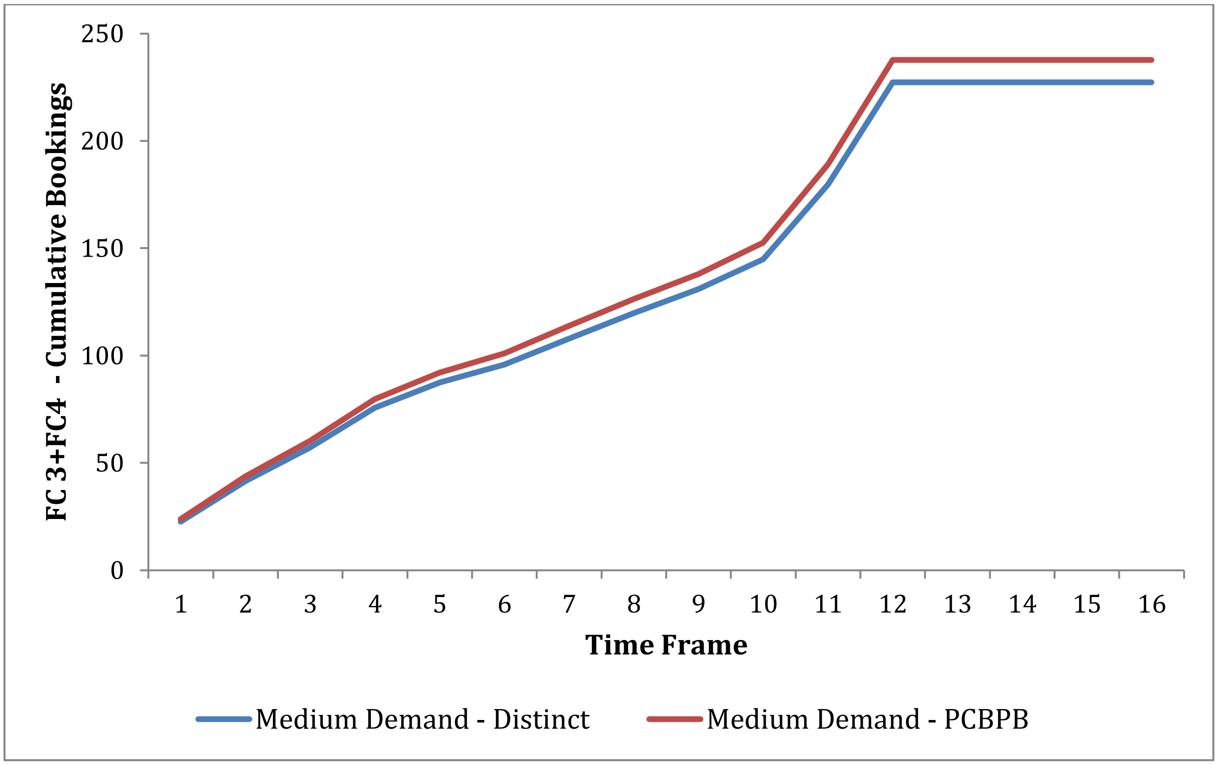

Figure 5-29: Cumulative Bookings FC3 + FC4 Medium Demand Level

The next group of fare classes to be analyzed is the group of high economy fare classes: FC5 and FC6. As presented in Figure 5-23, most of the revenue gains obtained by sharing premium cabin capacity with economy fare classes comes from revenue increases in FC5 with some additional gains in FC6. This result is not surprising considering the characteristics of the algorithm of Premium Cabin Bid Price Capacity Sharing. Since capacity is shared only when economy cabin is sold out, and this is expected to happen only in the last time frames, only the economy fare classes that are still open at such late time frames would be able to take advantage of the premium cabin capacity made available. Considering the advance purchase (AP) requirements in place for the economy cabin fare classes, booking requests for FC5 and FC6 are the only ones being accepted within the week before the flight.

Despite the increase in the economy cabin bid prices, when considering together the increase in the FC5 booking forecasts and the additional capacity made available, it is expected that the closure rate for FC5 should be lower when premium cabin capacity is shared than when the capacity of each cabin was managed distinctly. Effectively, the reduction in the closure rates for FC5 and FC6 beginning on time frame 11 is considerable and keeps increasing as the departure date is closer, as presented in Figure 5-30; for example, the closure rate in the last time frame decreases from 29.2% to 15.9% at the medium demand level. As a result of the reduction of the closure rates, the chances of acceptance of a booking request for these fare classes increase, and therefore more bookings are recorded. Combining FC5 and FC6 bookings, Figure 5-31 presents the increase in bookings in such fare classes, and how such increase is mostly concentrated in the last time frames of the booking period.

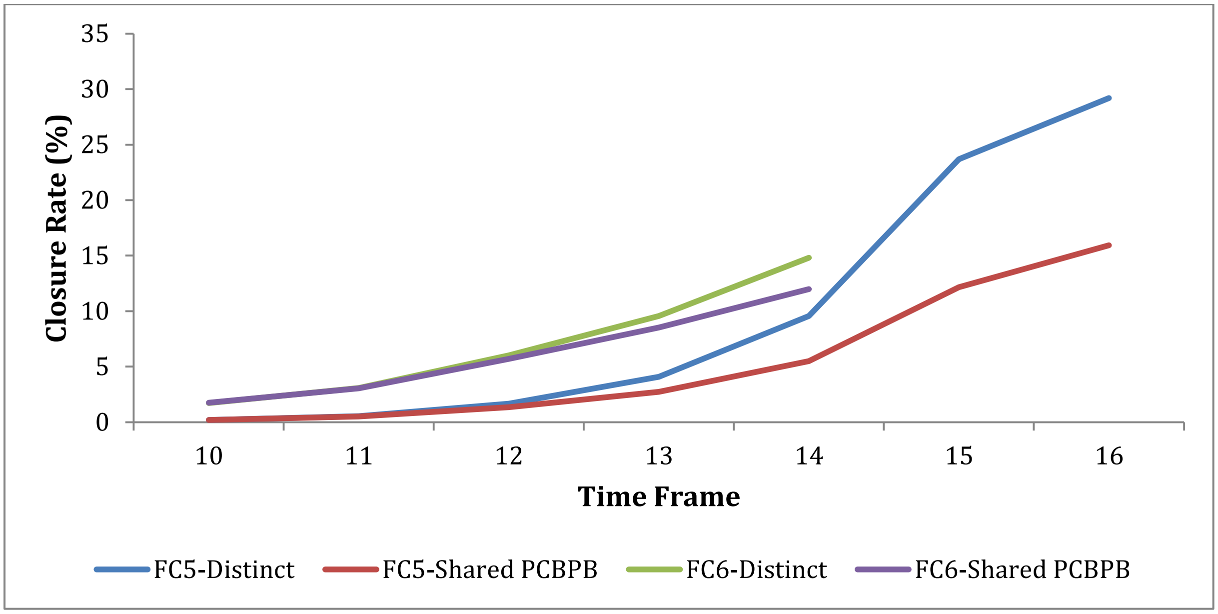

Figure 5-30: Closure Rates Fare Classes 5 and 6

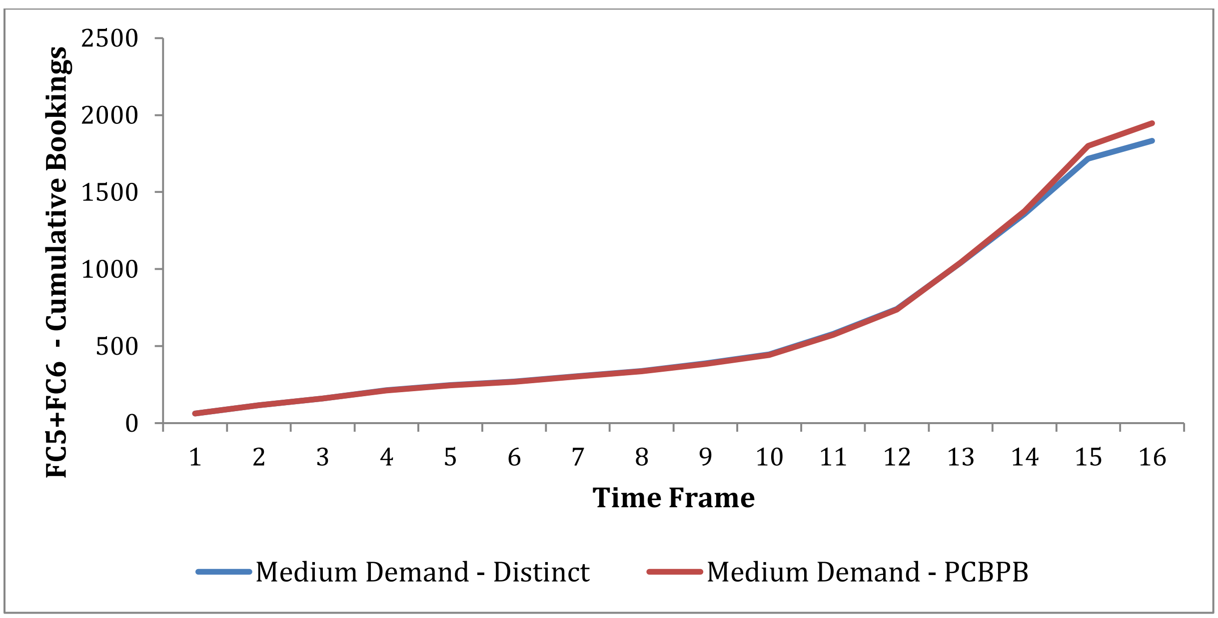

Figure 5-31: Cumulative Bookings FC5 + FC6 Medium Demand Level

The last group of fare classes is composed by FC7, FC8, FC9 and FC10; the closure rates for each of these fare classes are provided in Figure 5-32. Different from what has been observed in the cases of FC5 and FC6, this group of low economy fare classes share in common that the closure rate increases slightly when the premium cabin capacity is shared with economy fare classes using Method 1 of capacity sharing. This is a result of the higher economy cabin bid price observed for the sharing scenario. However, the differences between the closure rates are small enough to have a relevant impact in the revenue of AL1, with the sole exception of FC9 and FC10 in the high demand scenario; in those specific exceptions, the increase in the closure rate explains the revenue losses presented in Figure 5-23. The low variation in the closure rates explains the negligible difference between the total bookings in the low fare classes when premium capacity is shared using Method 1 and when capacity of each cabin is managed distinctly.

Finally, after examining the behavior of each fare class, a final insight associated with this premium cabin capacity sharing method is presented in Figure 5-33: it shows the proportional variation in total revenue during the booking period resulting from switching from distinct management of the cabins to sharing using the Premium Cabin Bid Price Capacity Sharing mechanism. In the initial timeframes, when many of the low fare economy classes are still open, the higher closure rates for such classes hurts the total revenue; however, as the end of the booking period approaches, the additional revenue captured from additional bookings in FC5 and FC6 produce enough revenue to compensate for the losses in FC1 and FC2 and generate a total revenue gain as a result of capacity sharing.

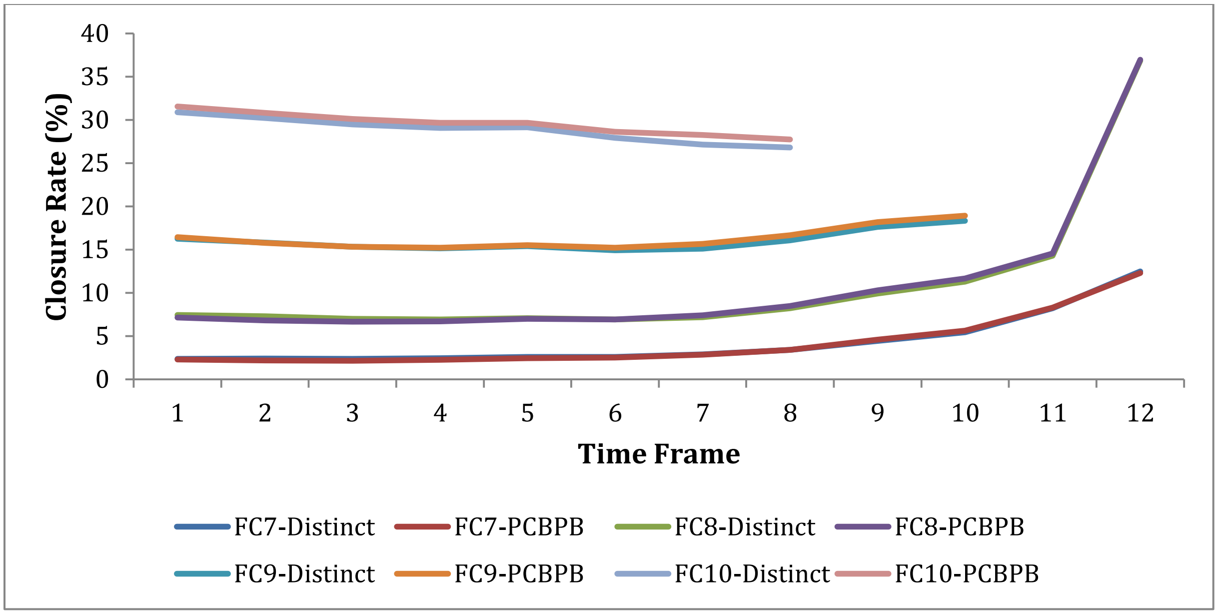

Figure 5-32: Closure Rates Fare Class 7 to 10 – Medium Demand

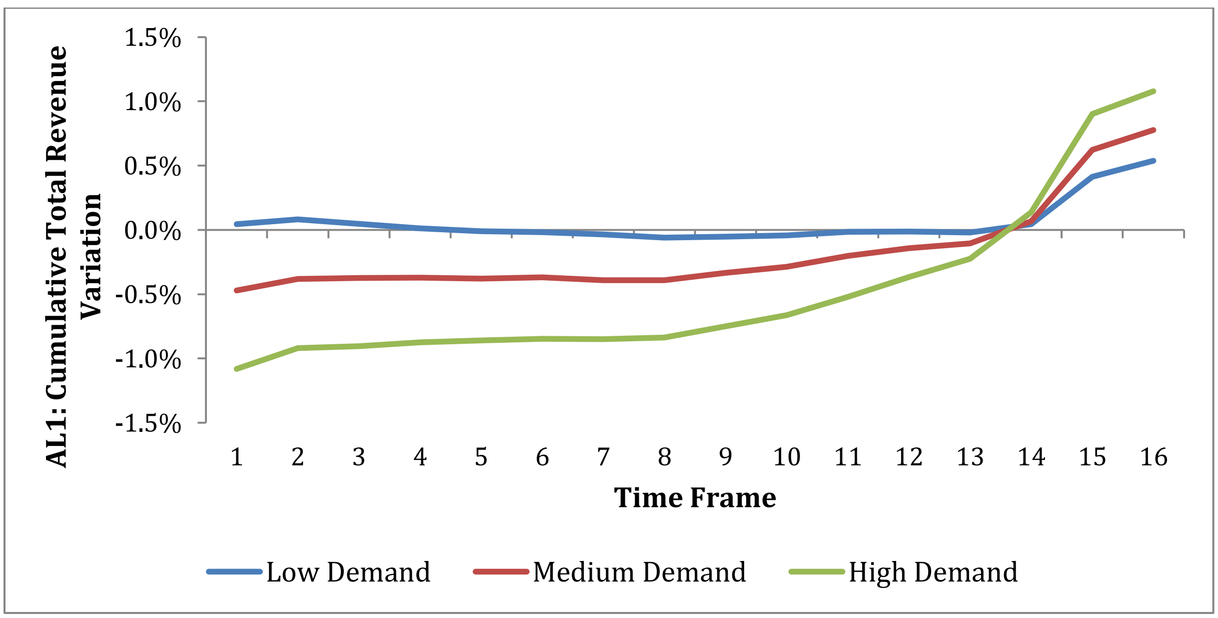

Figure 5-33: Cumulative Total Revenue Variation

5.2.2 Method 2: Joint Cabin Bid Price Capacity Sharing

In order to evaluate the performance of the Joint Cabin Bid Price Capacity Sharing heuristic (hereinafter referred to as "Method 2" or "JCBP") the RM systems of AL1 are set as described in Table 5-9. The only modification with respect to the OD RM Methods Baseline is that the mechanism implemented for sharing premium cabin capacity is the second one described in Section

3.5.2; that is, after economy cabin is sold out, premium cabin capacity is made available to premium and economy fare classes if the fare of such classes is higher than a joint bid price that considers the whole capacity of the aircraft cabin, ignoring any difference between premium and economy cabins (for connecting markets, the premium or economy fare must be higher than the sum of the joint cabin bid prices of the legs traversed by the ODF).

OD-Control Joint Cabin Bid Price Capacity Sharing
Network V1
AL1: Shared ProBP
Joint Cabin Bid Price Capacity Sharing (Method 2)
Path-based Std. Forecasting
AL2: Distinct ProBP – Path-based Std. Forecasting
AL3: Distinct EMSRb – Leg- Based Std. Forecasting
AL4: Distinct ProBP – Path-based Std. Forecasting

Table 5-9: Joint Cabin Bid Price Capacity Sharing - Simulation settings

The impact of the application of Method 2 by AL1 on the total revenue (premium plus economy) of each of the airlines operating in the network is presented in Figure 5-34 for three different demand levels. AL1 proportional revenue increases are between 0.58% and 1.16%, depending on the demand level (the higher the demand, the higher the proportional gain). Meanwhile, each of the other airlines competing in the market are negatively affected in their revenues by AL1's decision to make premium cabin capacity available for economy fare classes, as was the case with Method 1.

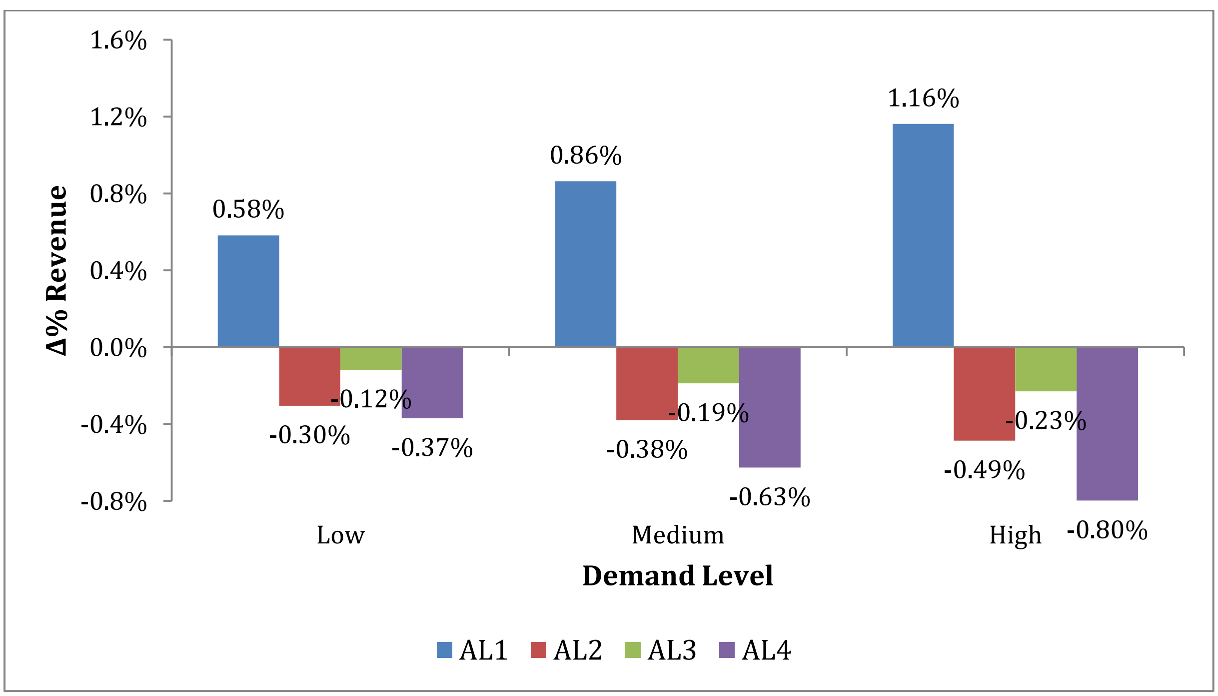

Figure 5-34: Method 2 - Total Revenue Proportional Variation by Airline

It should be noted as well that the revenue increases captured by AL1 with this Method 2 are slightly higher than those obtained when using Method 1, presented in the previous section; as a matter of fact, the difference in the total revenue captured of each of these methods is less than 0.1%. For example, although both mechanisms use very different approaches for sharing premium cabin capacity with economy fare classes, the Joint Cabin Bid Price capacity sharing methodology also generates proportional losses from premium fare classes (-4.11% to -9.76%) and gains from economy fare classes (1.65% to 3.70%), as shown in Figure 5-35. Again, since the economy fare classes' revenue represents most of the revenue of AL1, Figure 5-36 shows how the absolute gains in economy fare classes largely compensate the absolute revenue reductions in fare classes 1 to 4. Figure 5-37 illustrates that while the revenue captured from passengers flying in economy cabin (all of them with bookings in economy fare classes) increases, the combined revenue of the premium and economy fare class passengers flying in premium cabin decreases. Finally, Figure 5-38 shows how the revenue changes are concentrated in FC1, FC2, FC5 and FC10, just as observed as with the Premium Cabin Bid Price capacity sharing mechanism.

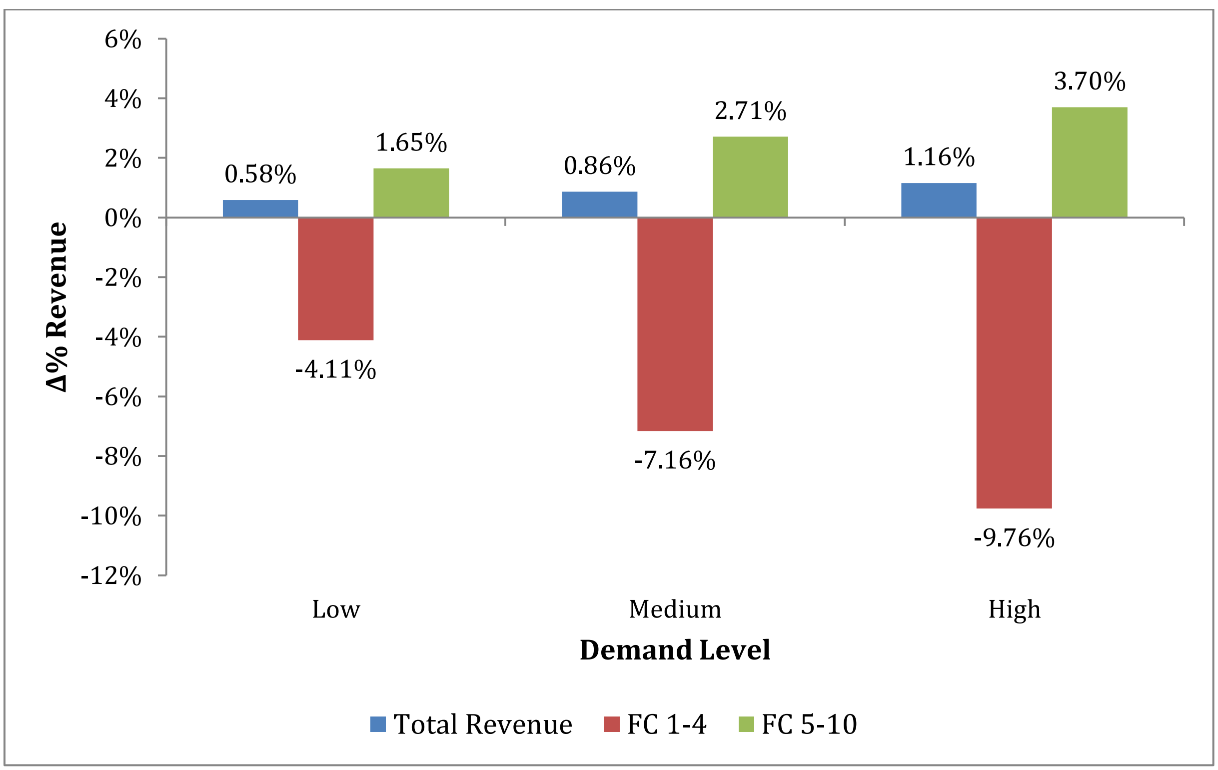

Figure 5-35: Method 2 - Proportional Revenue Gains by Fare Class Type

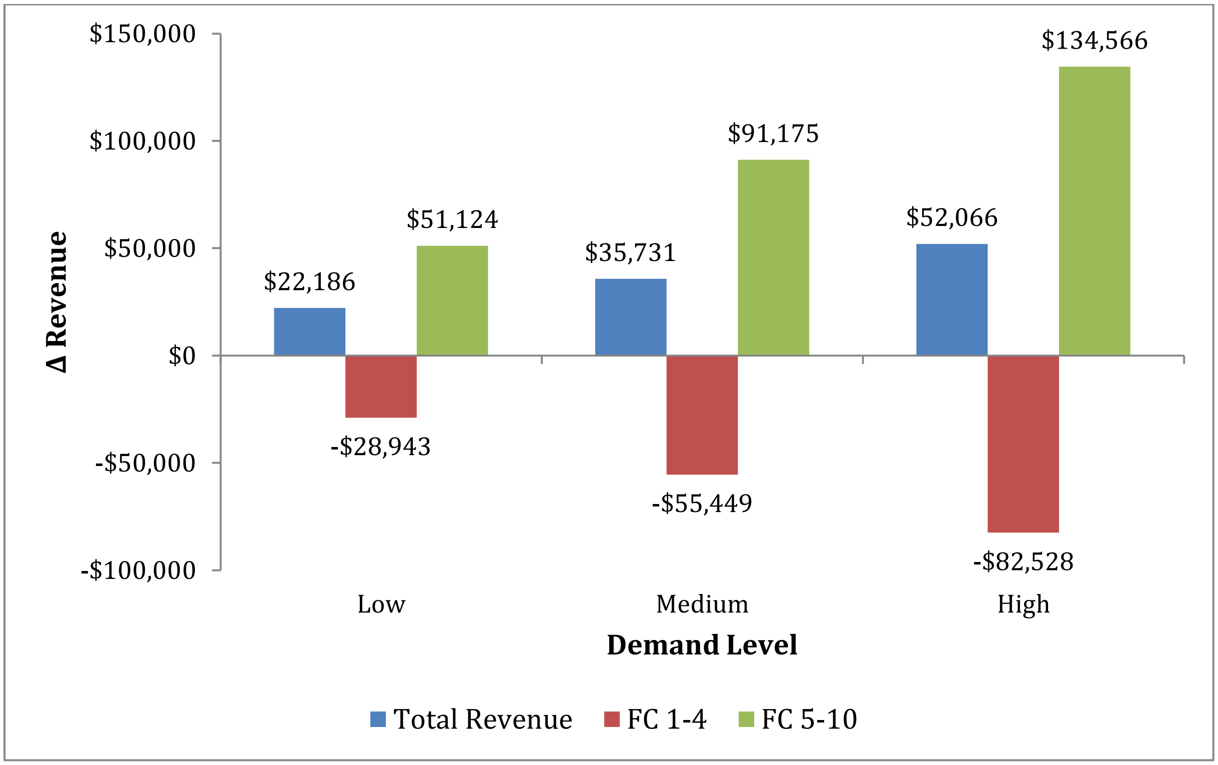

Figure 5-36: Method 2- Absolute Revenue Gains by Fare Class Type

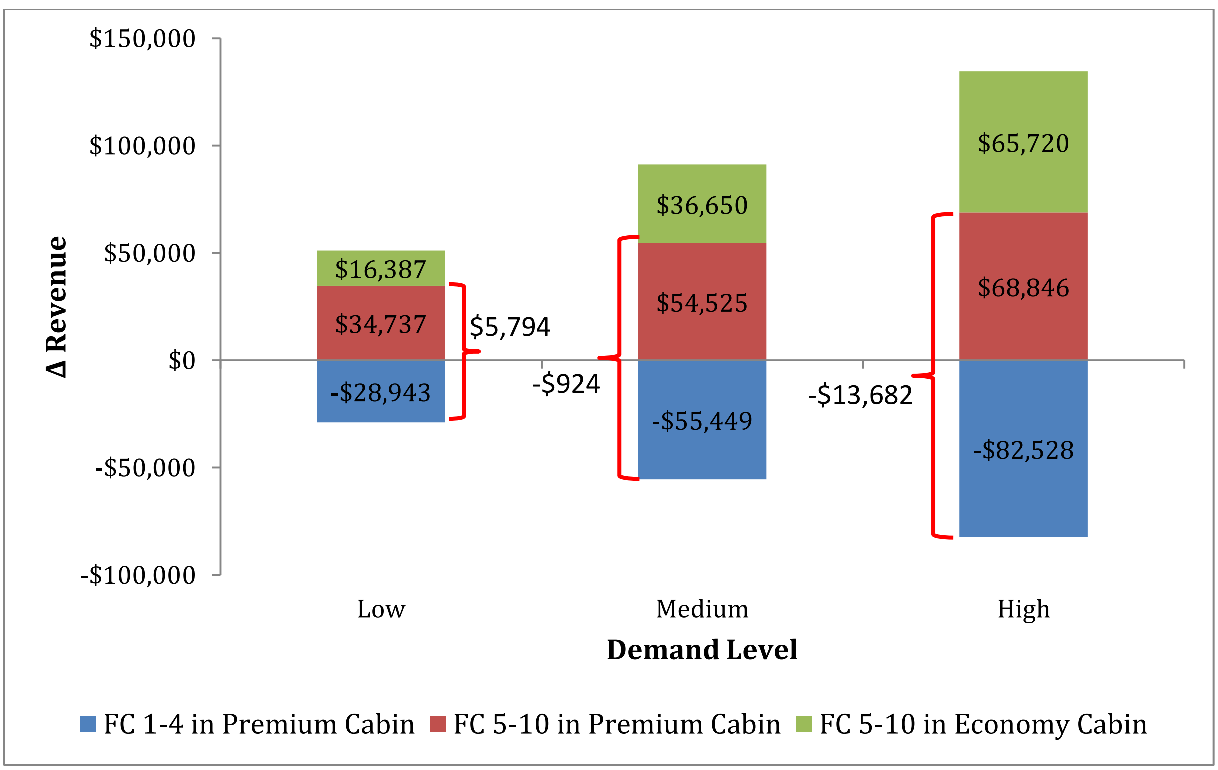

Figure 5-37: Method 2 - Absolute Revenue Variation by Fare Class Type and Cabin Accommodation

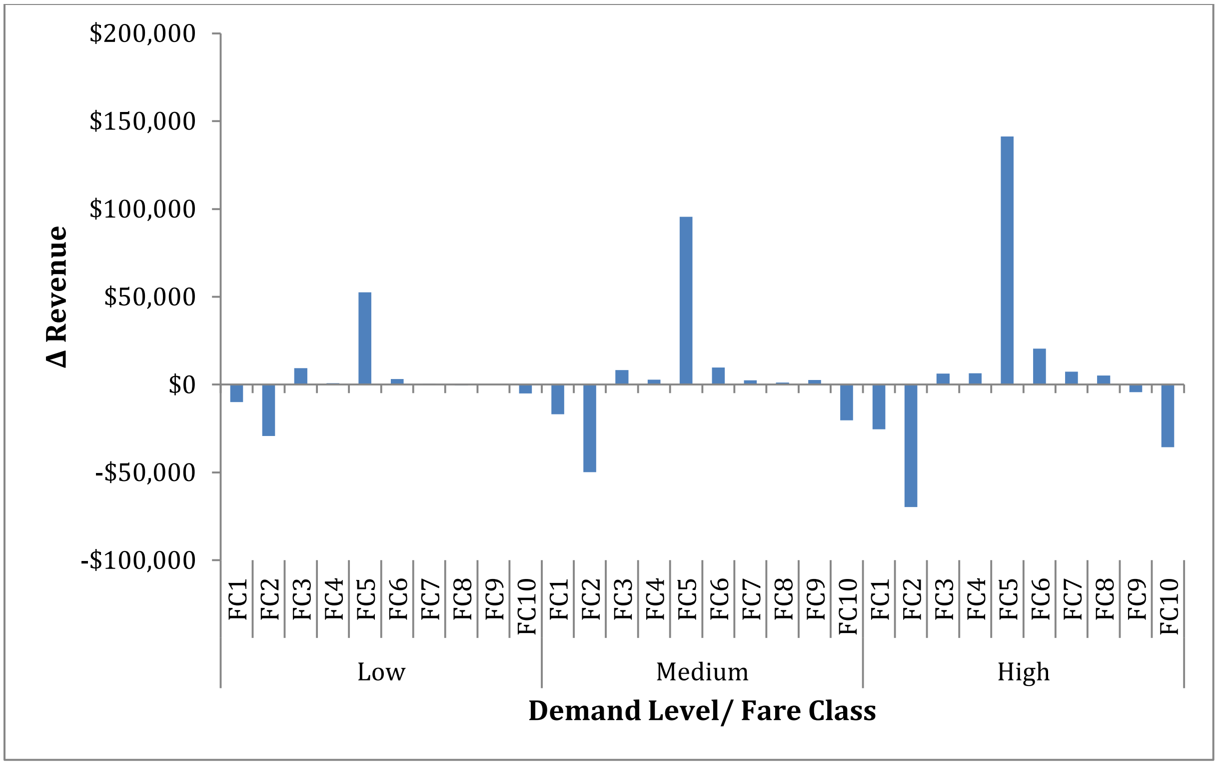

Figure 5-38: Method 2 - Revenue Variation by Fare Class and Demand Level

The similarity between the results of Method 1 and Method 2 suggests that the logic for analyzing both methods should be the same: on the one hand, the reduction in the number of bookings in premium fare classes results in a decrease of the forecasts for such fare classes and, consequently, lower premium cabin bid prices. On the other hand, bookings in economy fare classes increase as additional premium cabin capacity is made available; with more bookings in AL1's records, forecasts for economy fare classes are also increased and economy cabin bid prices are higher.

Confirming this approach, Figure 5-39 presents the average premium cabin bid prices for the scenarios where cabins are managed distinctly and where capacity is shared using Method 1 and Method 2, in the medium scenario. The first finding is that the average premium cabin bid prices are almost the same when using Method 1 and Method 2; this means that before economy cabins are sold out the performance of both methods should be very similar. In addition, the joint cabin bid price applicable for the corresponding demand level is included. The relevance of presenting this joint cabin bid price is that it is the bid price that is used for accepting booking requests after economy cabin is sold out; again, this is expected to happen late in the booking period. As presented in the figures, in the last time frames the average premium and joint cabin bid prices are relatively low (with respect to the fare of the premium fare classes available at that moment) and the difference between them are minor. However, a lower bid price controlling availability allows to increase slightly the booking requests accepted and to reduce the losses in premium fare classes' revenue with respect to the baseline.

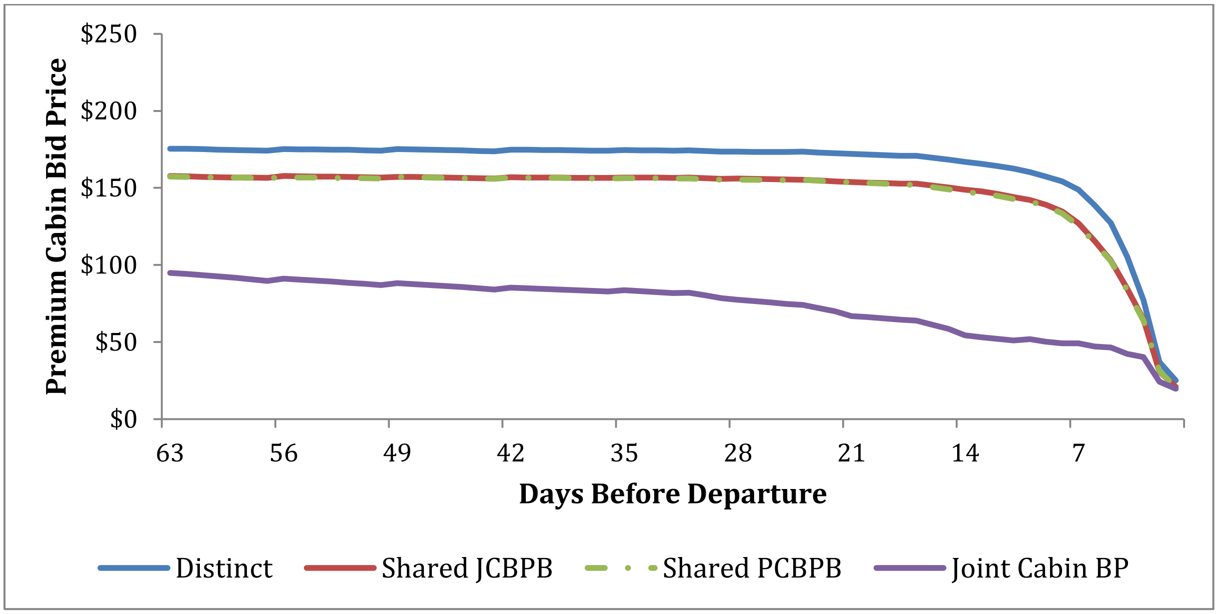

Figure 5-39: Premium and Joint Cabin Bid Prices – Medium Demand

A similar approach is used to analyze the results for economy fare classes: Figure 5-40 shows that for the three demand levels simulated the economy cabin bid prices for Method 1 (PCBPB) and Method 2 (JCBPB) are almost the same throughout the booking period. Therefore, the same interpretation given for Method 1 applies when comparing Method 2 with the baseline: higher economy bid prices protect more seats for high economy fare classes (FC5 and FC6). As a consequence of this, the applicable closure rates for such classes is lower in the last time frames, allowing for more high value bookings in the last days of the booking period.

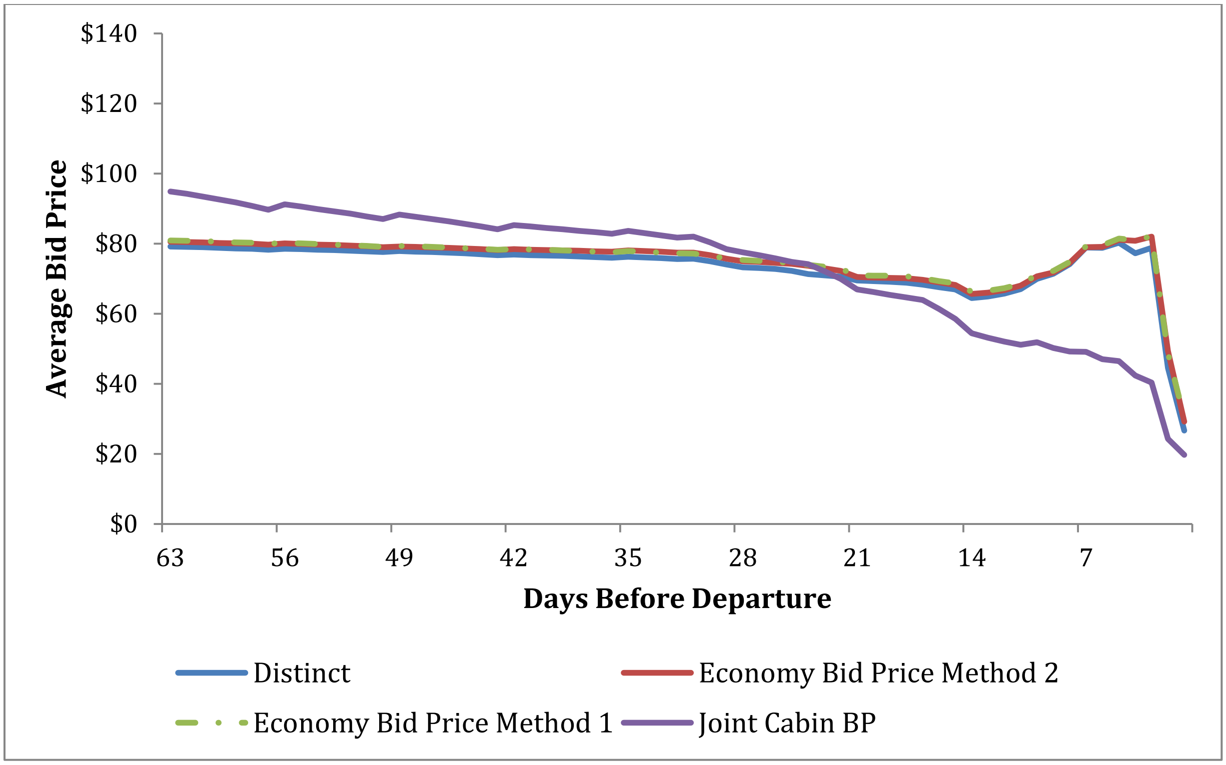

Figure 5-40: Economy and Joint Cabin Bid Prices – Medium Demand

The difference between Method 1 and Method 2 for economy fare classes is that once economy cabin is full, booking requests are evaluated using premium cabin bid price in the former and joint cabin bid price in the latter. Therefore, the same logic discussed for premium fare classes can be used to analyze Method 2: in the last time frames the average premium and joint cabin bid prices are relatively low (also with respect to the fare of the highest economy fare classes available at that moment) and the differences between them are minor. However, a slightly lower bid price is still able to stimulate few more bookings, increasing slightly the revenue gains from economy fare classes.

5.2.3 Effective Bid Price Capacity Sharing

The last premium cabin capacity sharing mechanism explored for bid price control RM is Effective Bid Price Capacity Sharing (hereinafter "Method 3" or "EBP"): the main settings for the simulations used to evaluate this method are presented in Table 5-10. As described in Chapter 3, this mechanism also calculates a joint cabin bid price for each leg assuming that the total aircraft capacity is not divided in cabins. Such joint cabin bid price is used as the "effective premium cabin bid price" if it is higher than the premium cabin bid price applicable for a leg. On the other hand, the joint cabin bid price is used as the “effective economy cabin bid price” instead of the economy cabin bid price if the former is lower than the latter. It should be noted that it is not necessary to have economy cabin sold out in order to allow for premium cabin capacity sharing, so it could happen since the beginning of the booking period; this characteristic is critical for understanding the results that will be presented in this section.

OD-Control Premium Cabin Bid Price Capacity Sharing
Network V1
AL1: Shared ProBP
Premium Cabin Bid Price Capacity Sharing (Method 3)
Path-based Std. Forecasting
AL2: Distinct ProBP – Path-based Std. Forecasting
AL3: Distinct EMSRb – Leg- Based Std. Forecasting
AL4: Distinct ProBP – Path-based Std. Forecasting

Table 5-10: Effective Bid Price Capacity Sharing - Simulation settings

The results of this method are considerably different from what has been observed from the other sharing mechanisms presented previously. As illustrated in Figure 5-41, a negligible increase in AL1’s 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. Compared to the results obtained with the other OD-control premium cabin capacity sharing mechanisms explored in this thesis, in which total revenue gains increase, it is easy to conclude that Method 3 is a less beneficial heuristic for sharing premium cabin capacity. A surprising finding is that, despite the losses of AL1, none of the other airlines competing in the market is able to improve its total revenue; in some way, this result suggests that the decision of Airline 1 is leaving the whole industry worse off.

As in the other mechanisms, the total revenue variation consists on 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%), as presented in Figure 5-42. In this case, however, 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 5-43).

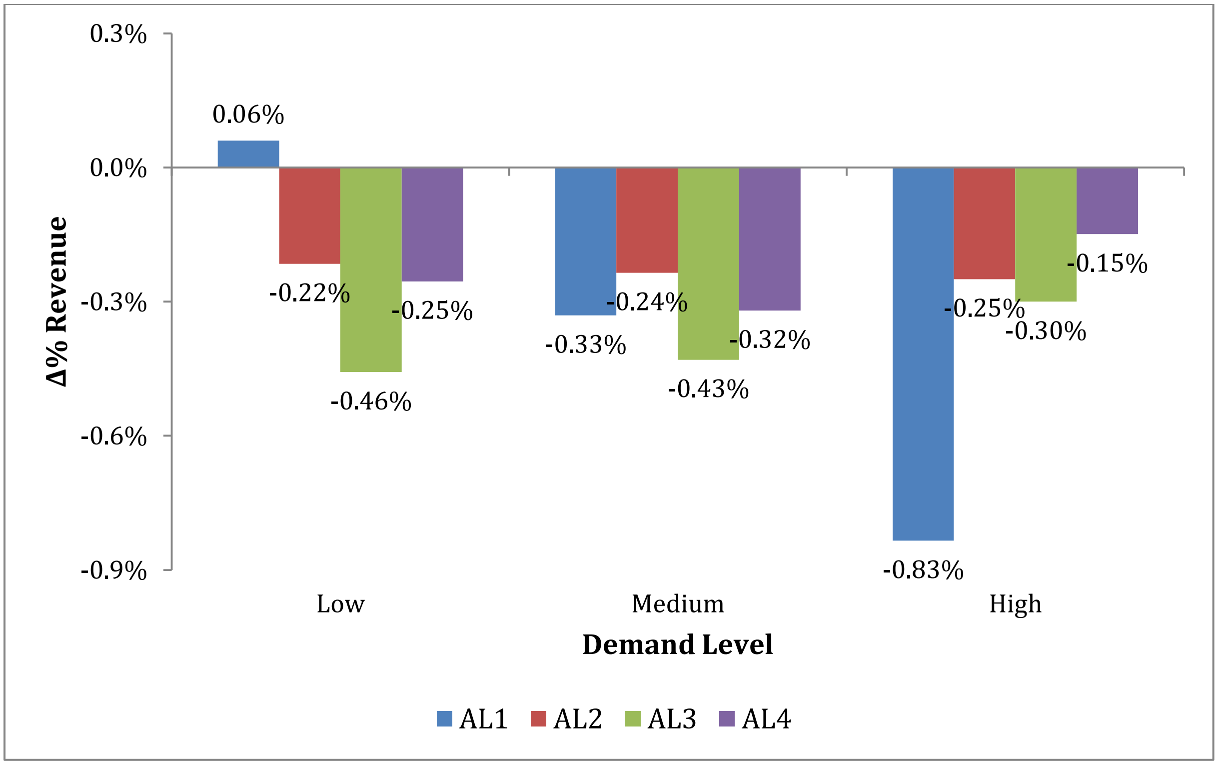

Figure 5-41: Method 3 - Total Revenue Proportional Variation by Airline

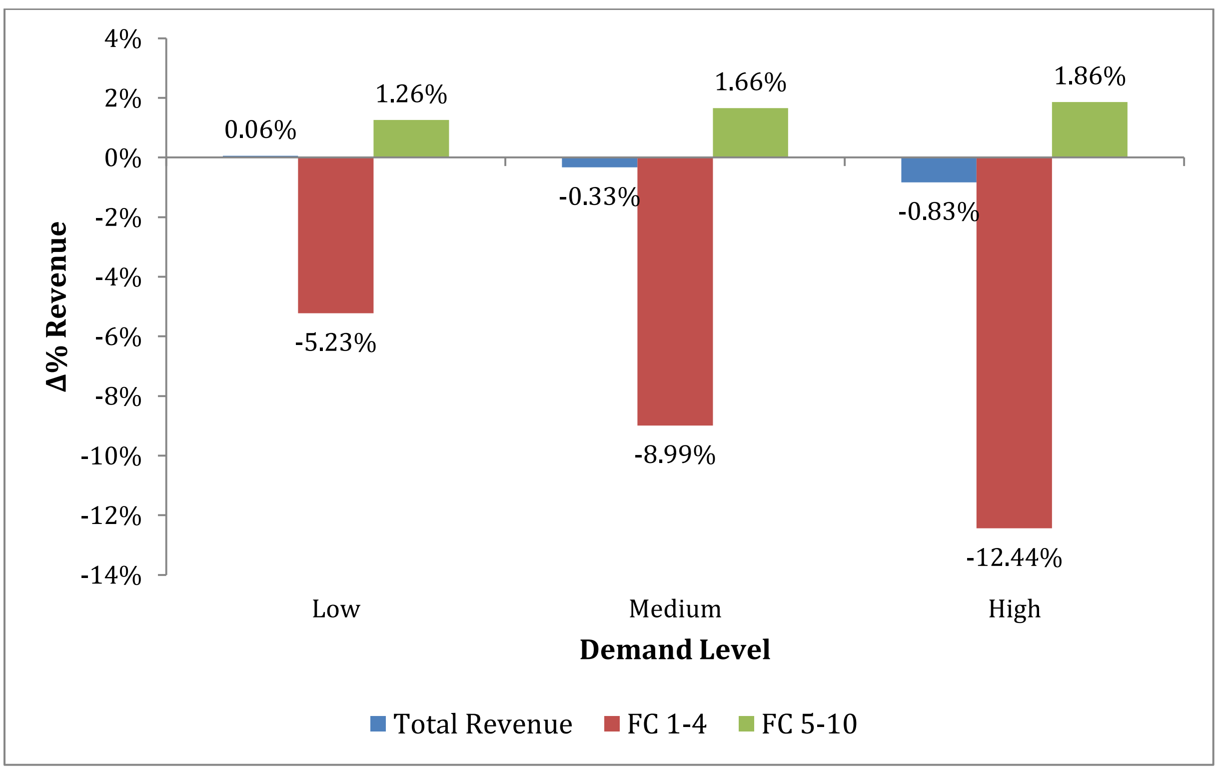

Figure 5-42: Method 3 - Proportional Revenue Gains by Fare Class Type

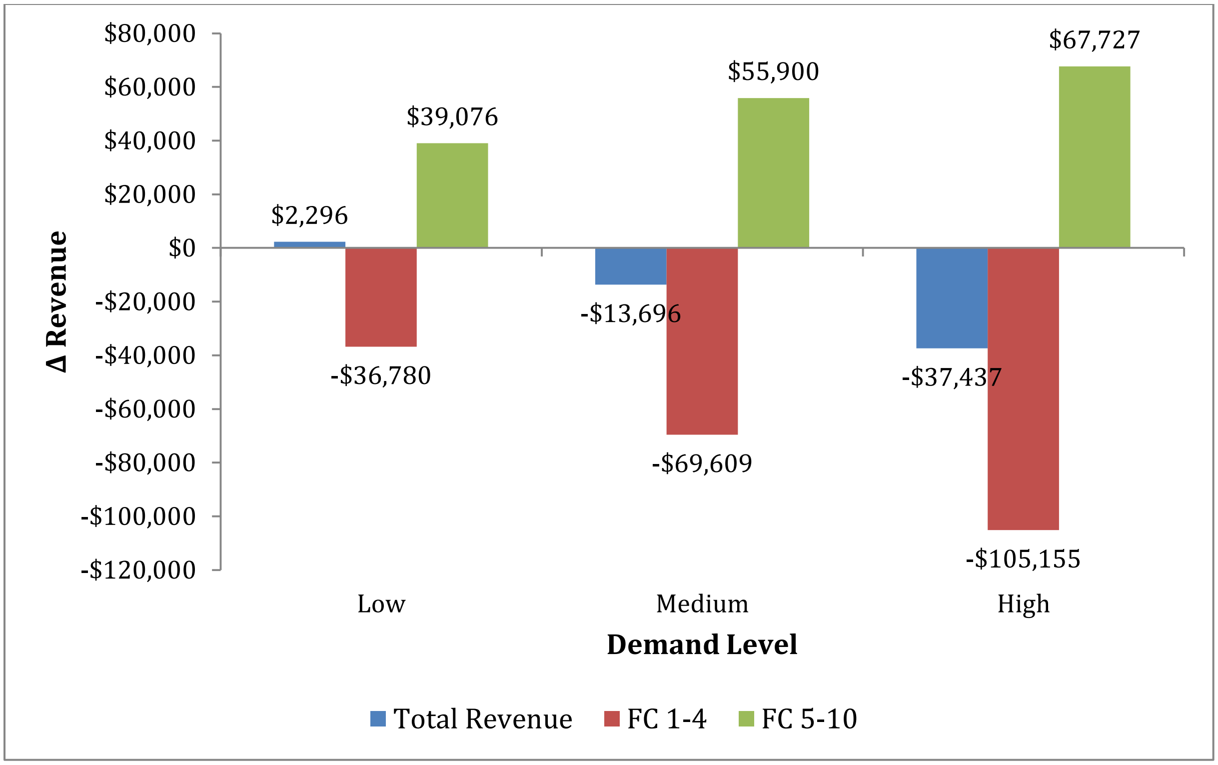

Figure 5-43: Method 3 - Absolute Revenue Gains by Fare Class Type

The results by cabin of Method 3 are very different to what has been found for the other two OD-control mechanisms presented previously in this chapter: specifically, EBP methodology captures a significant amount of revenue from passengers that booked in economy fare classes by allowing them to fly in 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. On the other hand, the revenue from passengers flying in economy cabin has a dramatic decrease that exceeds the premium cabin revenue gains, as presented in Figure 5-44. Once again, the simulations lead to different results compared to what would have been intuitively expected when premium cabin capacity sharing is implemented: Table 5-11 summarizes the differences between the expected outcomes and the results of the simulations.

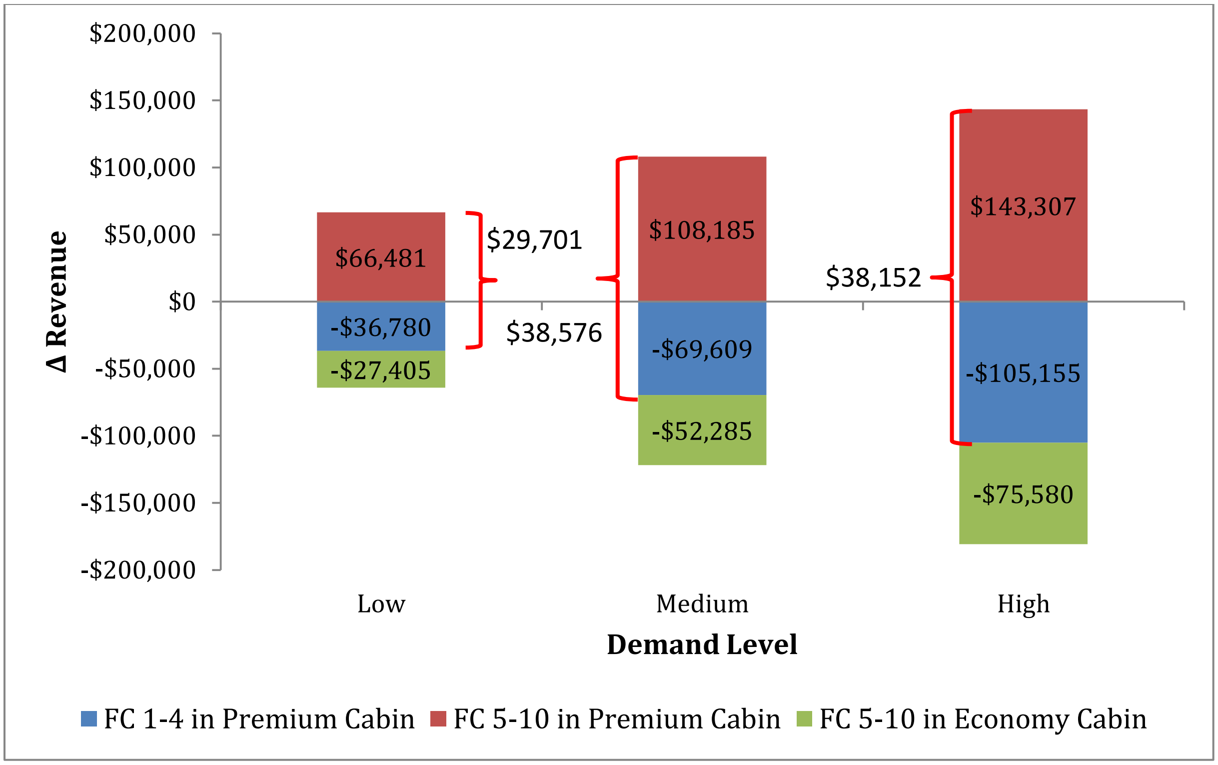

Figure 5-44: Method 3 - Absolute Revenue Variation by Fare Class Type and Cabin Accommodation

Concept	Expected	Simulations
Total Revenue (premium + economy)	Increases	Increases slightly/Decreases
FC 1-4 Revenue	Unchanged	Decreases
FC 5-10 Revenue	Increases	Increases
Premium Cabin Revenue	Increases	Increases
Economy Cabin Revenue	Unchanged	Decreases

Table 5-11: Comparison of Expected and Observed Results - EBP Capacity Sharing Method

A detailed examination of the revenue variations by fare classes gives some additional insights about the performance of the Effective Bid Price premium cabin capacity sharing method. As shown in Figure 5-45, the main changes are the losses in FC1 and FC2 and the gains in FC5 and FC10. The exchange of revenue losses in the highest premium fare classes for gains in the highest economy fare class has also been observed when Method 1 and Method 2 are used; however, this exchange is no longer beneficial for the airline since the losses in FC1 and FC2 dramatically increase when the demand increases while the incremental gains in FC5 revenue do not grow enough to compensate for such losses. It should be note also that FC10 revenue also increases considerably;

although these gains are not enough to compensate for the losses in the premium fare classes, it is useful to illustrate that when Method 3 is used, the effects of capacity sharing are observed even in the initial stages of the booking period.

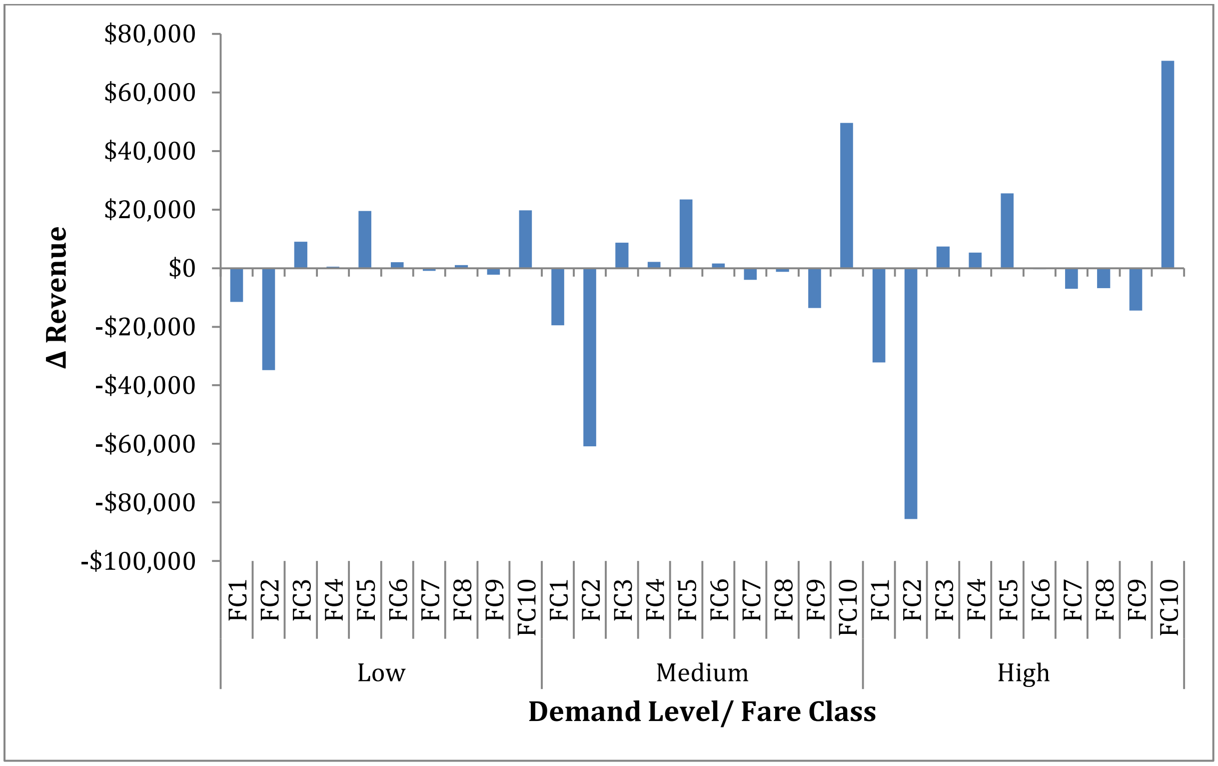

Figure 5-45: Method 3 - Revenue Variation by Fare Class and Demand Level

Again, as a consequence of the observation of less bookings because of having premium cabin capacity shared with economy fare classes, the forecasts for FC1 and FC2 end up being lower when compared to the baseline scenario, where capacities of each cabin are managed distinctly (Table 5-12). Comparing these numbers with the ones presented for Method 1 (Table 5-8), it can be identified that the reduction in the forecasts is higher with Method 3, mostly for the high demand scenario. In the case of FC5 a similar level of increase in the forecasts is observed for all demand levels; however, these increases are lower than the ones presented for Method 1.

Fare Class	Demand Level	Forecast Baseline	Forecast Method 1	Variation
FC1	Low	92.8	82.8	-10.8%
	Medium	123.7	105.2	-14.9%
	High	170.4	118.7	-30.4%
FC2	Low	487.0	393.1	-19.3%
	Medium	645.4	452.1	-30.0%
	High	891.8	536.0	-39.9%
FC5	Low	831.8	848.5	2.0%
	Medium	940.0	962.2	2.4%
	High	1061.2	1095.1	3.2%

Table 5-12: Method 3 - Proportional Variation of Forecasts

As described in previous sections, the impact of the changes in the forecasts is initially observed in the bid prices; in the case of the premium cabin, lower forecasts for premium fare classes result in lower premium cabin bid prices, as illustrated in Figure 5-46 for the medium demand scenario. Since joint cabin bid prices are of particular interest because of the design of the heuristic, the average joint cabin bid prices are also included in these figures. Method 3 determines that booking requests for premium fare classes (FC 1-4) are accepted only if the applicable fare is higher than the maximum value between the premium cabin bid price and the joint cabin bid price. Therefore, the fare of a premium fare class would have to be at least as high as it would have had to be if compared only with the premium cabin bid price in order to have a booking request accepted. Low average joint bid prices with respect to the average premium cabin bid prices would suggest that, in most cases, the premium cabin bid price is acting as the effective bid price for premium fare classes, so the result in such fare classes should not be as different as if only the premium cabin bid price was being used; on the other hand, the higher the average joint bid prices values are (with respect to the average premium cabin bid prices), the higher the likelihood of having a joint cabin bid price acting as the effective premium cabin bid price, resulting in a reduction of the bookings. As a matter of fact, for the medium and high demand scenarios the average joint bid price is higher than the average premium cabin bid price in the last days before departure.

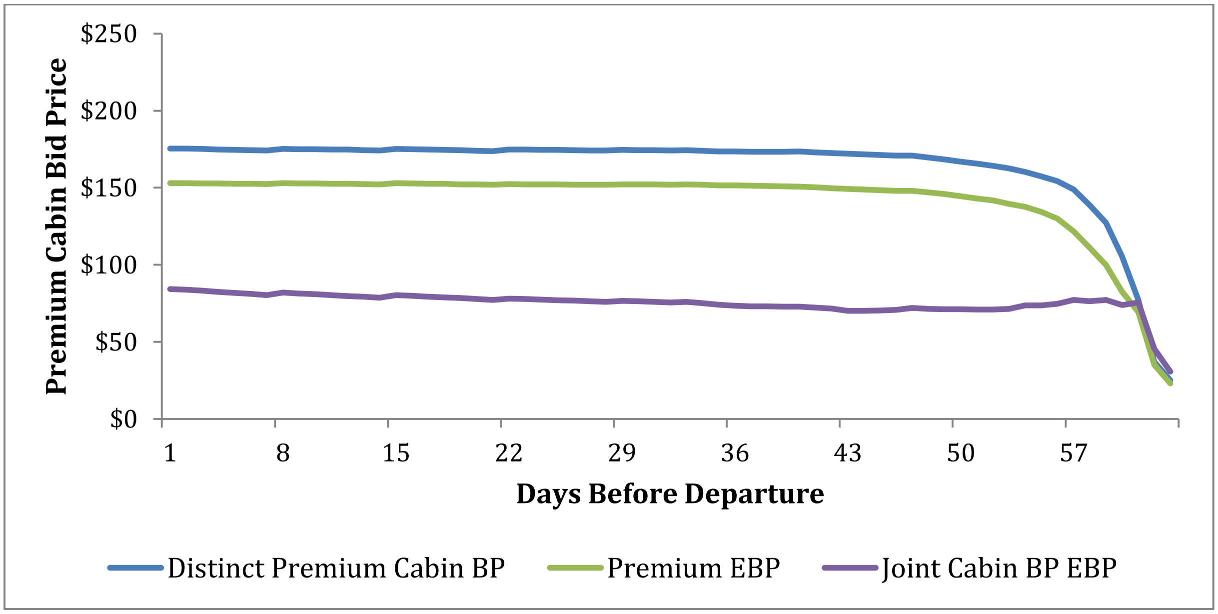

Figure 5-46: Premium and Joint Cabin Bid Prices – Medium Demand

The consequence of using higher bid prices for accepting (or rejecting) booking requests of premium fare classes is that the closure rates of such fare classes increases, as evidenced in Figure 5-47 for FC1 in the medium demand scenario. For the purposes of comparison, such figure includes as well the applicable closure rate with Method 1 (PCPB): closure rate with Method 3 is substantially higher, and explains the increase in the FC1 losses relative to Method 1.

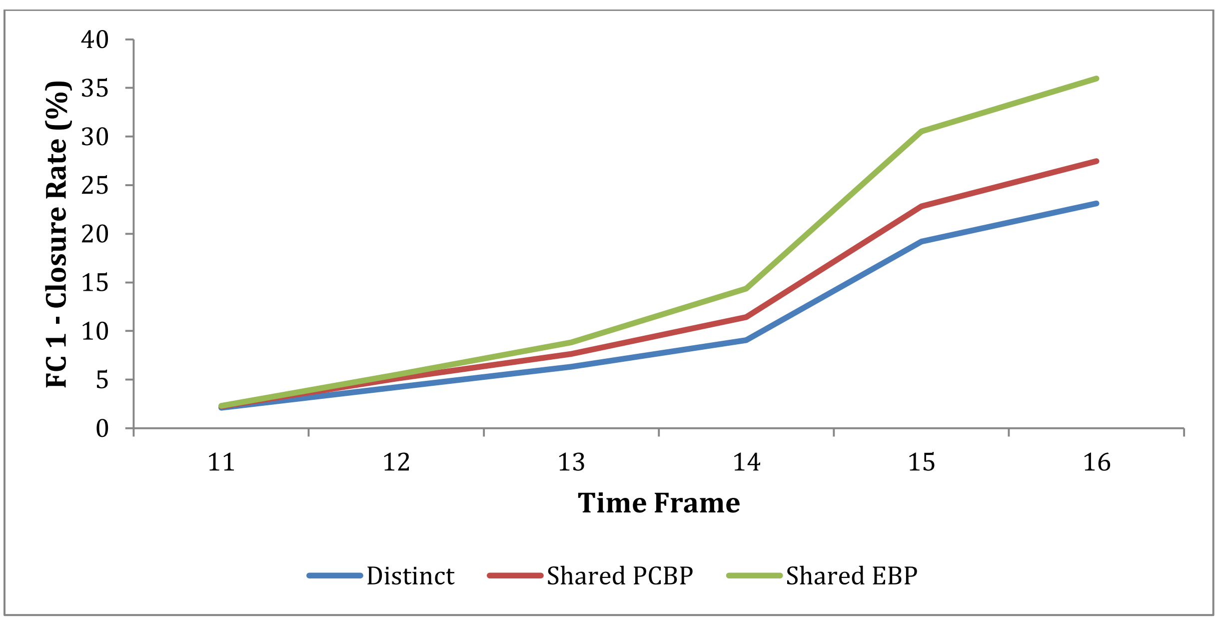

Figure 5-47: Closure Rate Fare Class 1 - Method 3 - Medium Demand

The high closure rates observed for the high premium fare classes with Method 3 help to understand the reduction in FC 1 and FC2 bookings observed in the last time frames of the booking process, as illustrated in Figure 5-48.

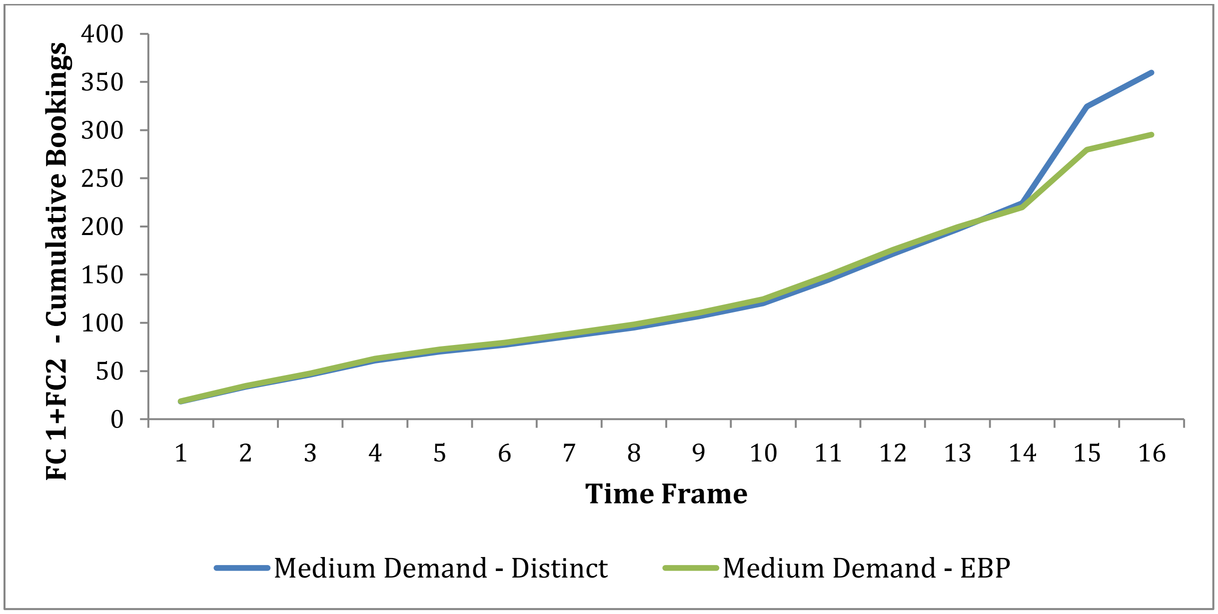

Figure 5-48: Cumulative Bookings FC1 + FC2 - Method 3 - Medium Demand

As discussed before, the heuristic used in Method 3 determines the "effective" economy cabin bid price by selecting the minimum between the economy cabin bid price and the joint cabin bid price. Therefore, booking requests that would have not been accepted if only the economy cabin bid price was being used could be accepted if the joint cabin bid price is lower; consequently, more low fare class bookings are expected to be accepted.

Figure 5-49 presents the average economy and joint cabin bid prices for Method 3 in the medium demand scenario as well as the average economy cabin bid price when cabins are managed distinctly. The behavior and interaction of these bid prices can be described in three stages: the first of is when the average economy bid prices for the EBP methodology are lower or equal to the one of the applicable for the distinct case. The second stage happens when the average economy and joint cabin bid price exceed the average distinct economy cabin bid price. Finally, in the last days of the booking period, either the joint or the economy EBP bid prices are lower than the distinct economy cabin bid prices. These stages matter because make it easy to understand the behavior of FC5 and FC10 bookings, already identified as the major sources of variation of Method 3 with respect to the baseline.

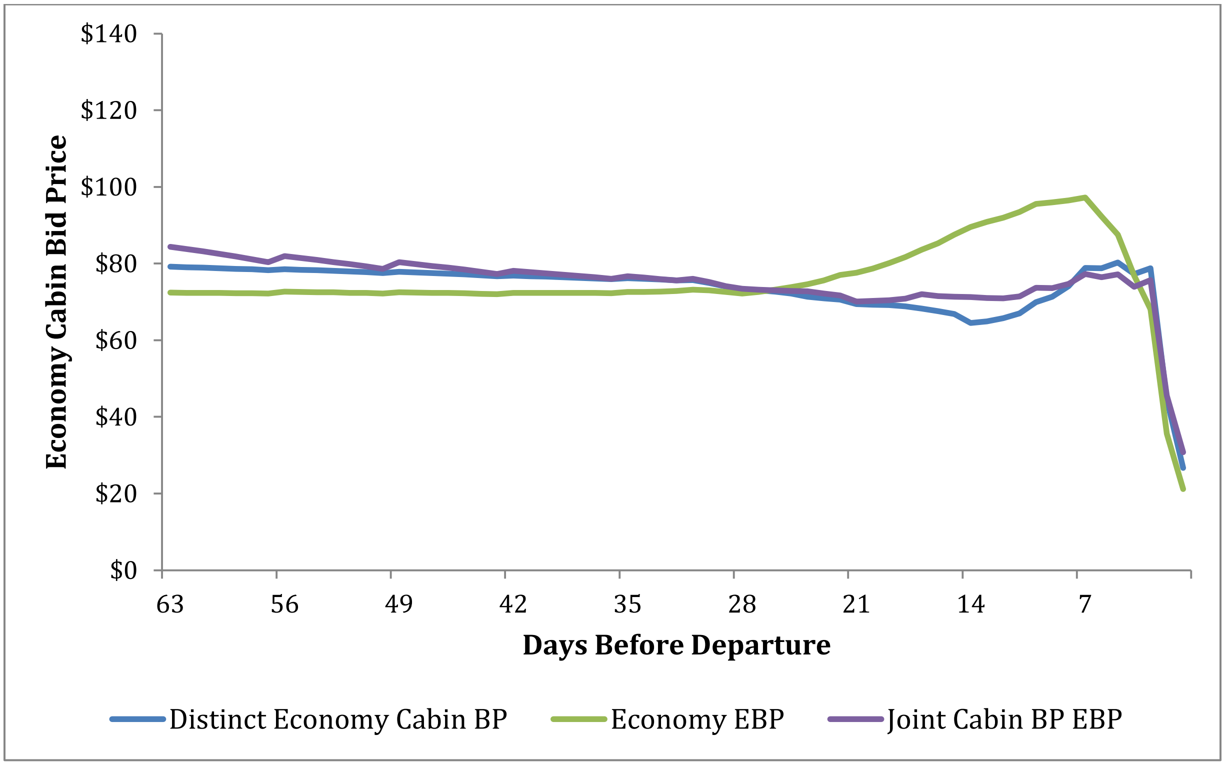

Figure 5-49: Economy and Joint Cabin Bid Prices – Medium Demand

The initial stage, where the average economy cabin bid price for Method 3 is lower than in the baseline, the likelihood of accepting a low fare booking request increases: this is reflected in a much lower closure rate for FC10, as presented in Figure 5-50 and, consequently, in increased FC10 bookings in the initial time frames of the booking period (Figure 5-51). On the other hand, the final stage more bookings in such fare class should be expected. This would suggest that in the initial time frames of the booking period, it is easier for a low fare (such as FC10) booking request to be accepted.

On the other hand, the final stage of the booking process is characterized by an average joint cabin bid price that is lower than the average economy cabin bid price of the baseline. Therefore, more booking requests should be accepted in the economy fare classes that are not closed because of advance purchase requirements (i.e. FC5 and FC6). Figure 5-52 and Figure 5-53 present, respectively, the closure rate and cumulative bookings curve for FC5, showing an increase in bookings that confirms this approach.

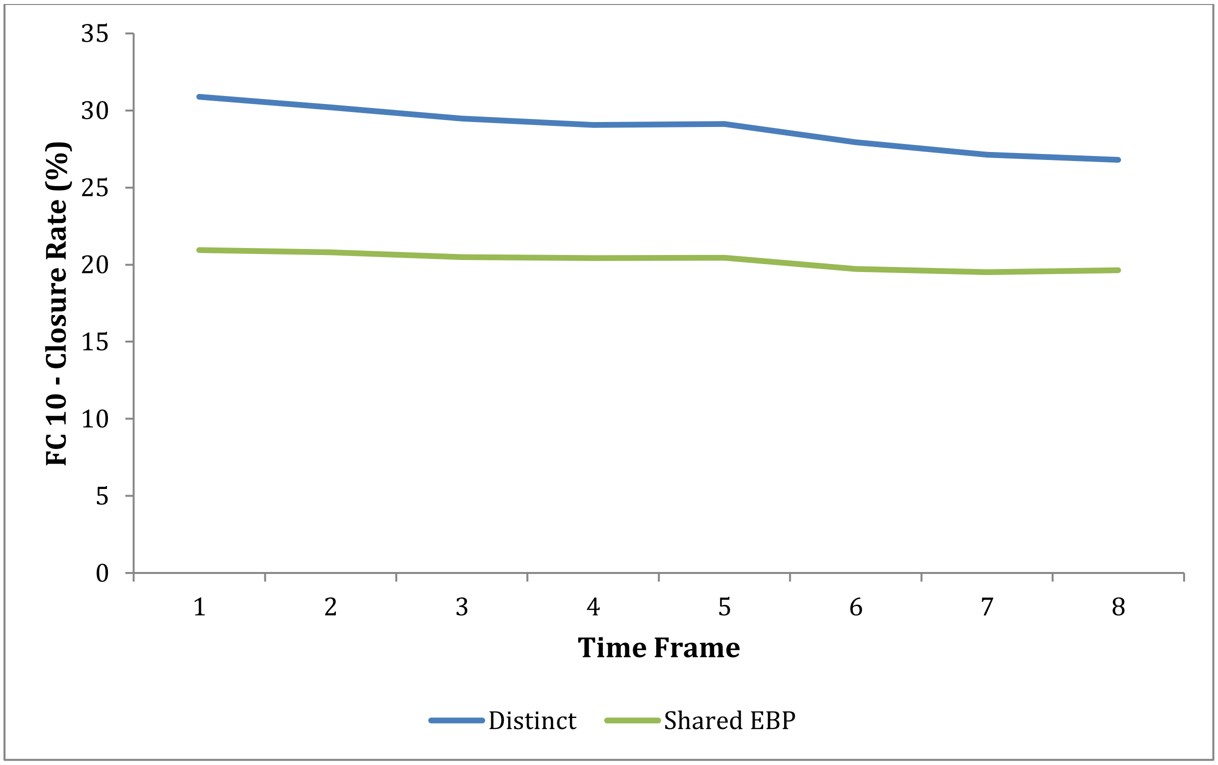

Figure 5-50: Closure Rate Fare Class 10 - Method 3 - Medium Demand

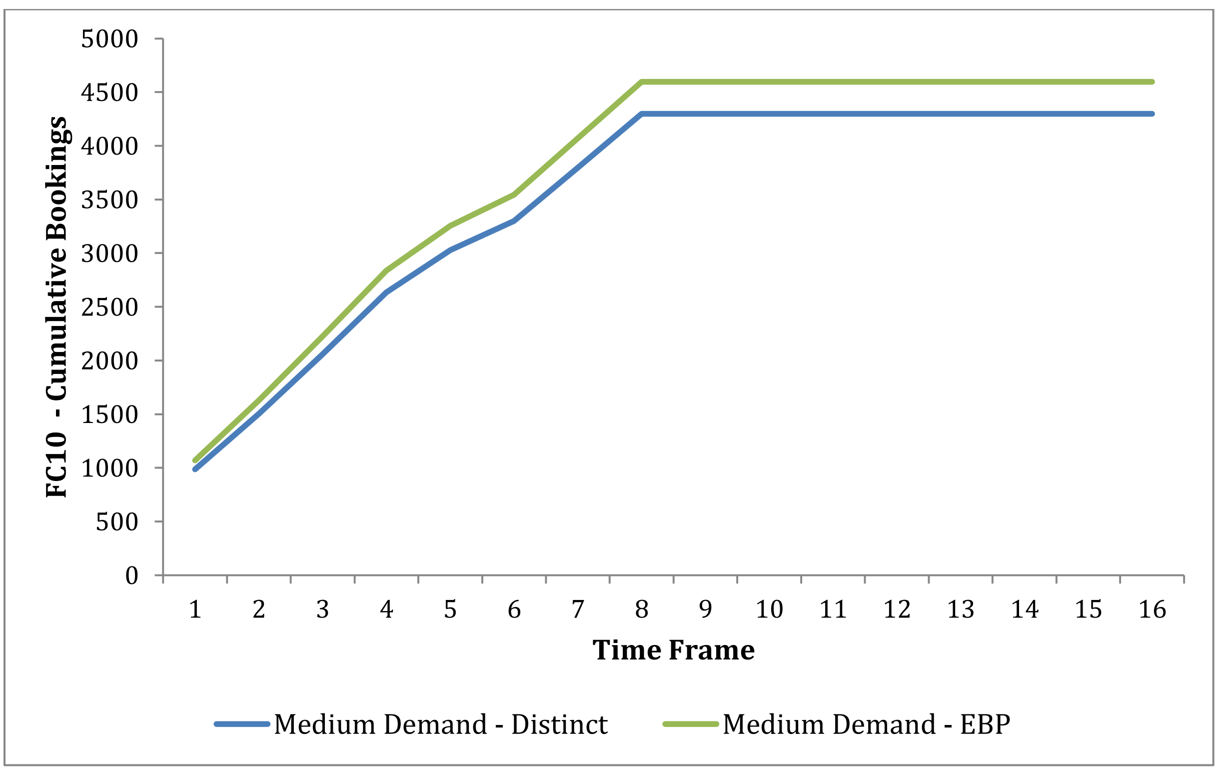

Figure 5-51: Cumulative Bookings FC10 - Method 3 - Medium Demand

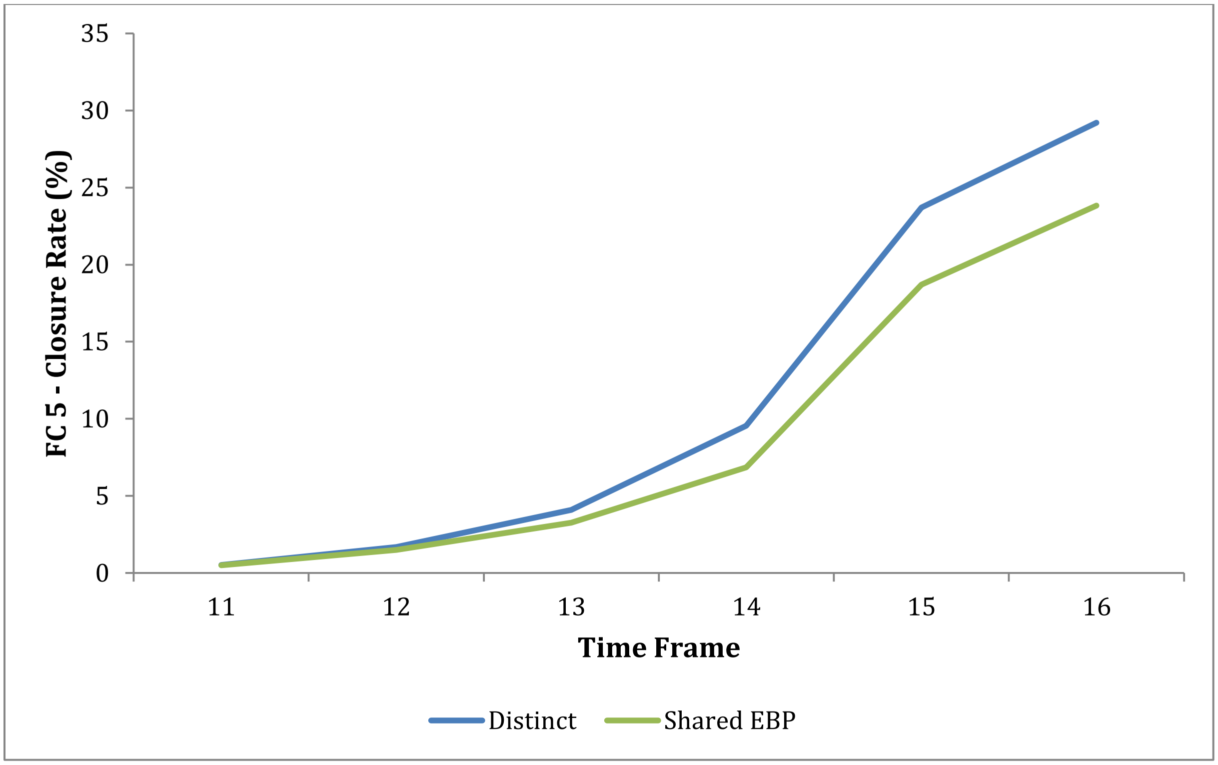

Figure 5-52: Closure Rate Fare Class 5 - Method 3 - Medium Demand

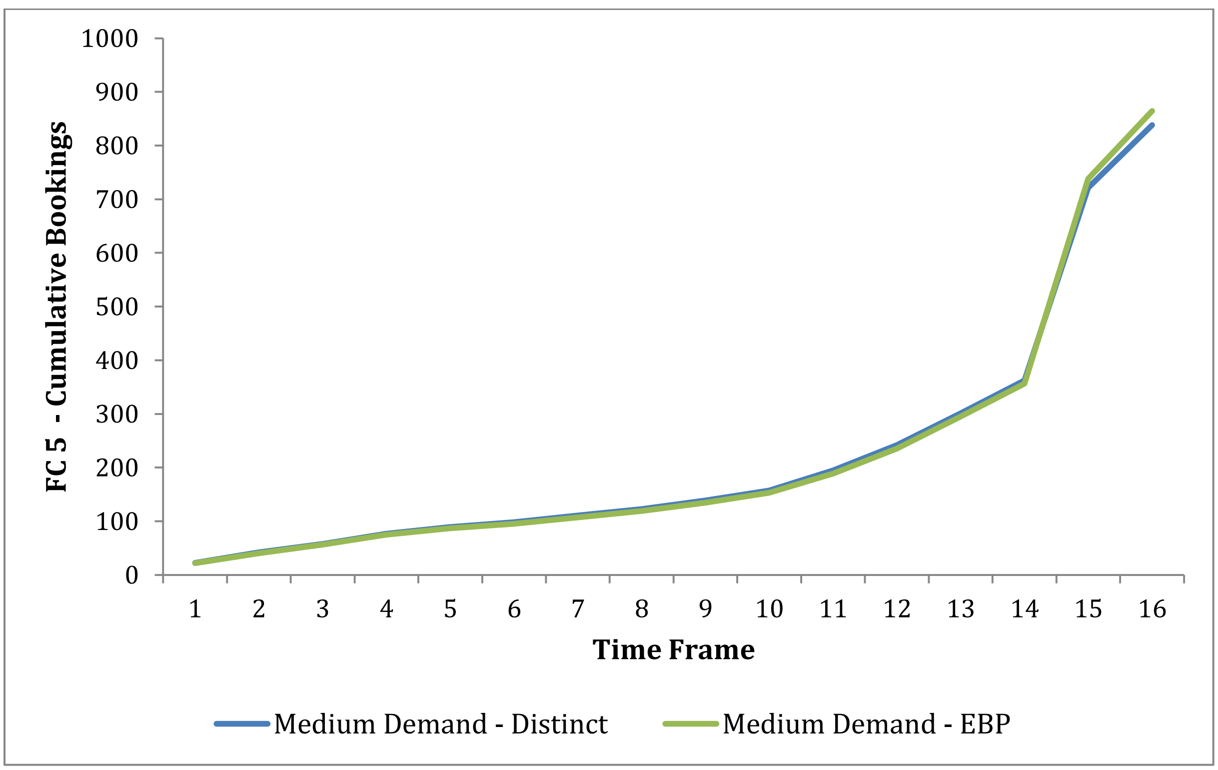

Figure 5-53: Cumulative Bookings FC5 - Method 3 - Medium Demand

Figure 5-54 presents the variation in cumulative total revenue at each demand level. In the initial time frames, the revenue captured using Method 3 is higher than in the baseline mostly because of the increase in FC10 bookings. However, in the last days before departure all the gains are lost because of the reduced availability of FC1 and FC2, despite the additional bookings recorded in FC5.

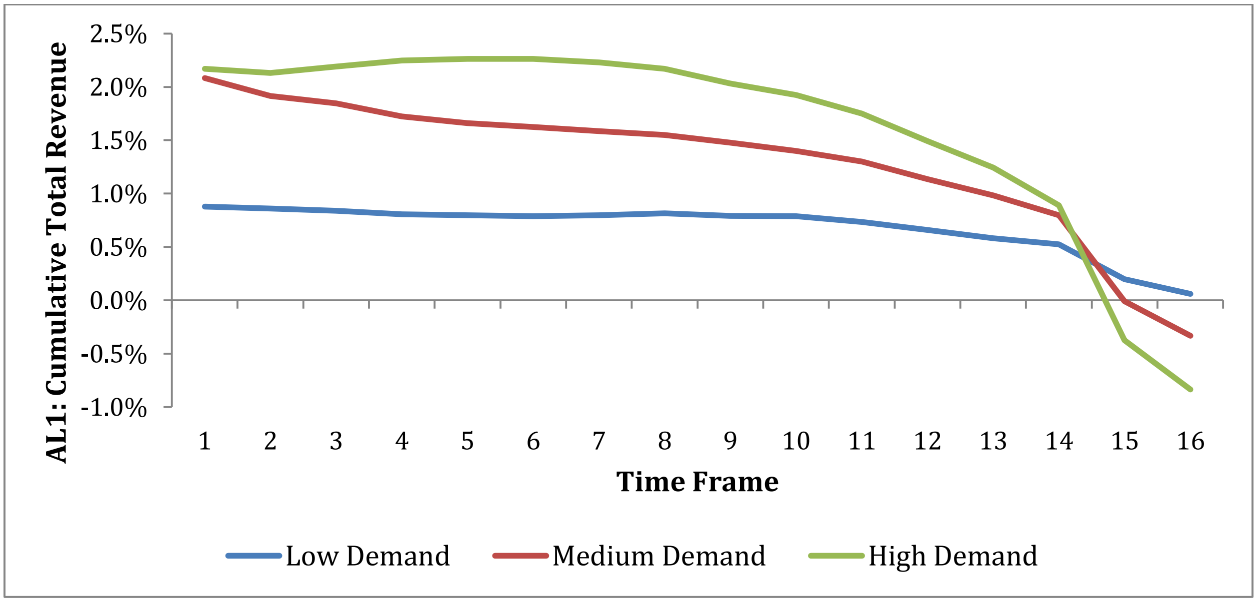

Figure 5-54: Cumulative Total Revenue Variation - Method 3

5.2.4 Protection Mechanisms for Premium Fare Classes – OD control methods

As discussed in the initial sections of this chapter, the motivation for implementing a premium cabin capacity sharing mechanism is increasing the revenue captured by accommodating additional passengers that are able to buy economy fare classes in premium cabin seats that otherwise would have been flown empty. Conceptually, these additional passengers should not replace or displace the passengers wishing to fly in premium cabin; however, the results of the simulations have shown that all the heuristics proposed result in a reduction in the revenue captured from FC's 1 to 4. Despite the increase in total revenue (premium + economy), airlines value greatly the passengers that buy high value premium fare classes and a reduction in that source of revenue is something that the management of airlines that operate dual cabin aircraft usually find problematic and not strategic in the long term.

Taking this into consideration, two mechanisms for protecting premium fare classes' revenue are proposed. The first of these consists in delaying premium cabin capacity sharing until a predetermined time frame of the booking process with the intention of guaranteeing that capacity is shared only with passengers that buy some of the high-value economy fare classes that are less constrained by AP requirements. The second scheme proposed consists of activating premium cabin capacity sharing on a leg only after a certain predefined economy cabin load factor has been reached; this approach is also useful to avoid sharing premium capacity throughout the booking process and guarantees that premium cabin capacity sharing is restricted only to flights with a considerable economy cabin load factor.

- **Time Frame Protection**

TFP is tested using the following predetermined time frames:

- Time Frame 10 (TFP=10): distinct until 14 days before departure
- Time Frame 12 (TFP=12): distinct until 7 days before departure
- Time Frame 14 (TFP=14): distinct until 3 days before departure

Each of the OD-Control RM Shared Capacity Schemes proposed herein is evaluated with each of the TFP parameters stated above. The first mechanism evaluated is the Premium Cabin Bid Price Capacity Sharing method (Method 1). The results of the tests at low, medium and high level are included as Figure 5-55 and details about the medium demand scenario are illustrated in Figure 5-56. Roughly speaking, the total revenue gains are kept at the same level as in the scenario without TFP at all demand levels; moreover, the changes in proportional gains in economy fare classes and losses in premium fare classes resulting from introducing TFP are negligible for TFP=10 and TFP=12, at all the demand levels. However, some progress towards the objective of reducing the revenue losses for FC's 1 to 4 is achieved with TFP=14: for example, losses are reduced from -7.31% to -6.94% in the medium demand scenario, compensating for a reduction in revenue gains in economy fare classes from 2.64% to 2.55% and keeping the proportional revenue gains stable at 0.78%.

Similar results are obtained from the application of TFP to Method 2: Figure 5-57 illustrates that the impact of TFP is negligible. Overall, the results for both methods confirm that the requirement to share only after economy cabin is sold out prevents the negative aspects associated with sharing premium cabin capacity too early in the booking process, as observed with Method 3.

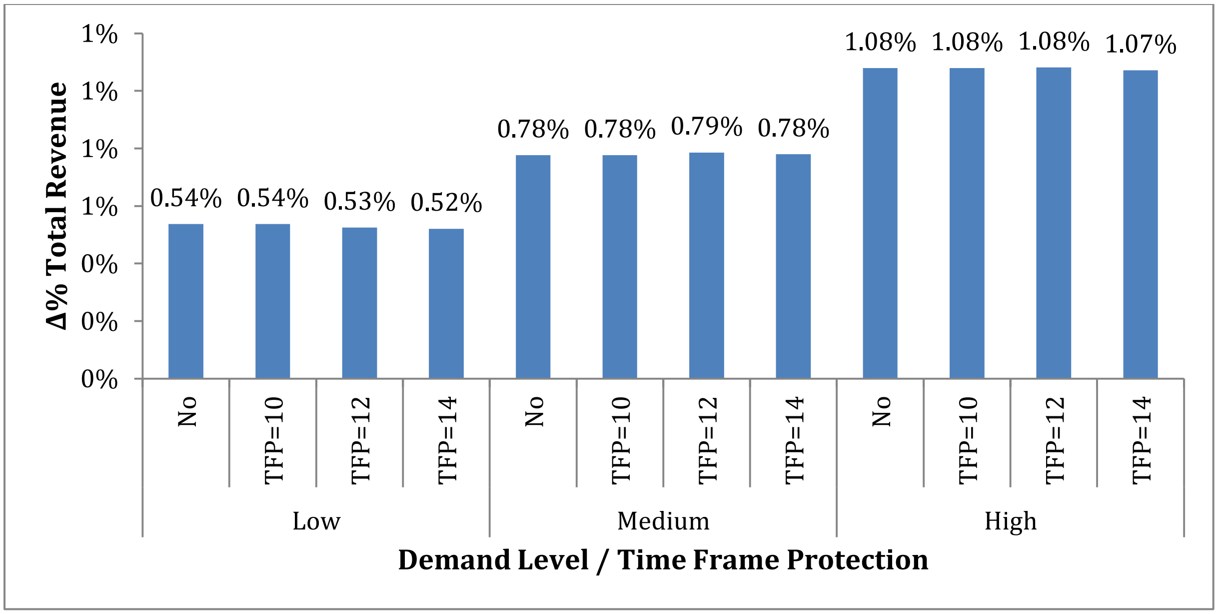

Figure 5-55: Time Frame Protection - Method 1

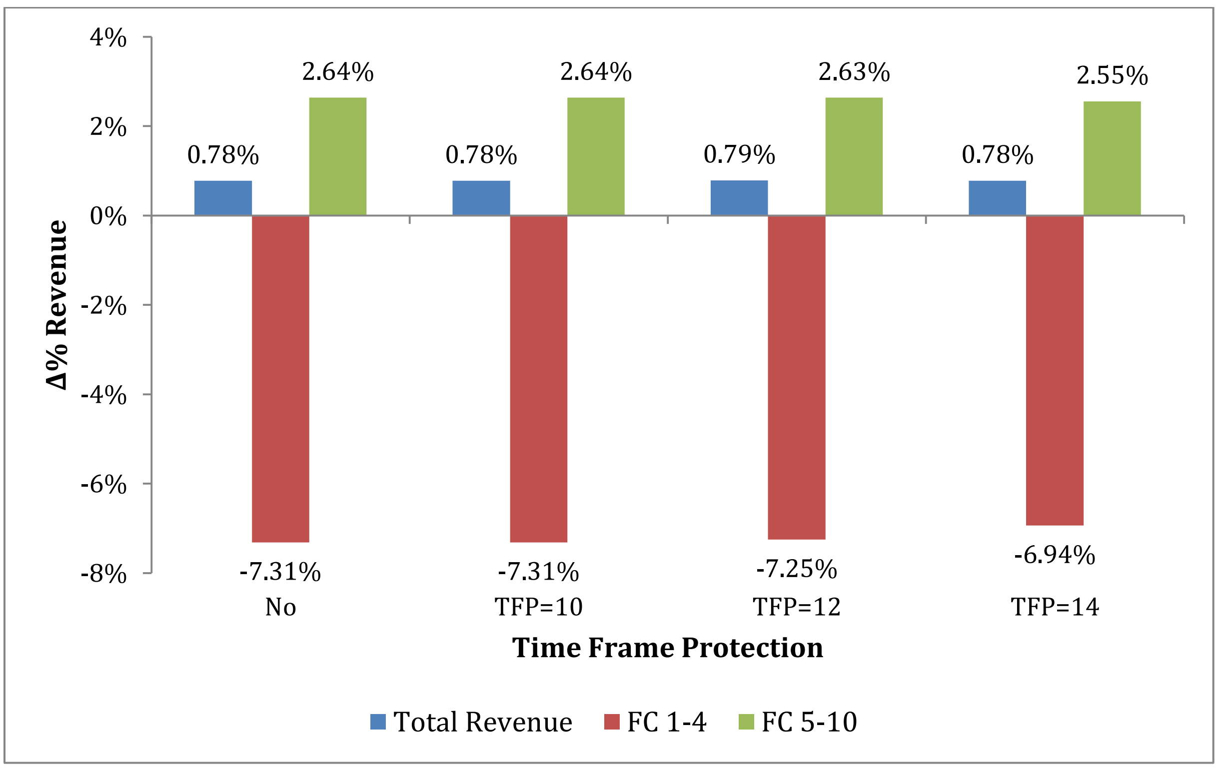

Figure 5-56: Time Frame Protection - Method 1 - Medium Demand

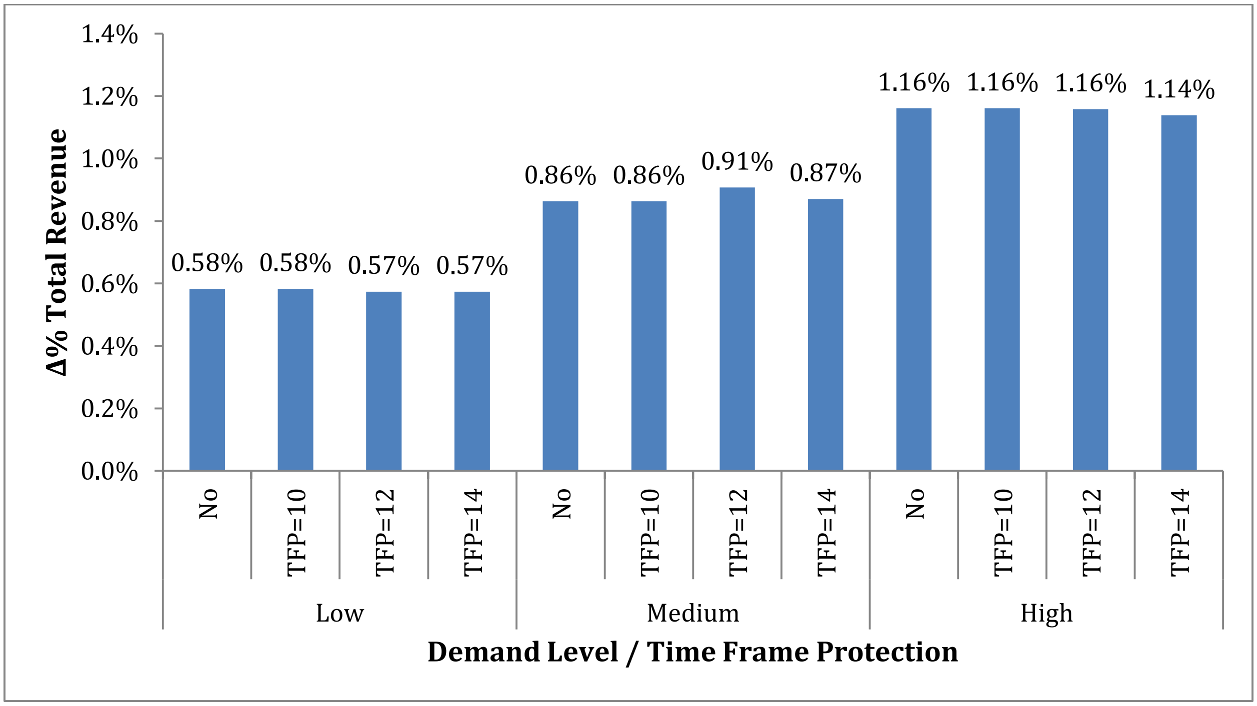

Figure 5-57: Total Revenue - Time Frame Protection - Method 2

The impact of TFP on the performance of the Effective Bid Price premium cabin capacity sharing method is much more important than what has been found for Method 1 and Method 2. While in the low demand scenario the gains increase from 0.06% to 0.58% (Figure 5-58), in the medium and high demand scenarios TFP turns losses into gains; as a matter of fact, although the proportional losses obtained in the high demand scenario without TFP are larger than in the medium demand scenario, the highest proportional revenue gains are achieved in the high demand scenario with TFP=14 (1.1%) (Figure 5-59 and Figure 5-60). 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; an example of this is provided in Figure 5-61 for the medium demand scenario.

Other interesting aspects of the effect of TFP on the performance of Method 3 can be found in the proportional variations for each category of fare classes. One of the findings is that this protection mechanism is successful in reducing the proportional losses in premium fare classes; for example, in the medium demand scenario such losses decrease from -8.99% to -8.47%. The other relevant insight is that the revenue captured from economy fare classes also increases with the introduction of TFP, despite this protection mechanism is supposed to represent a constraint for the additional capacity made available to this group.

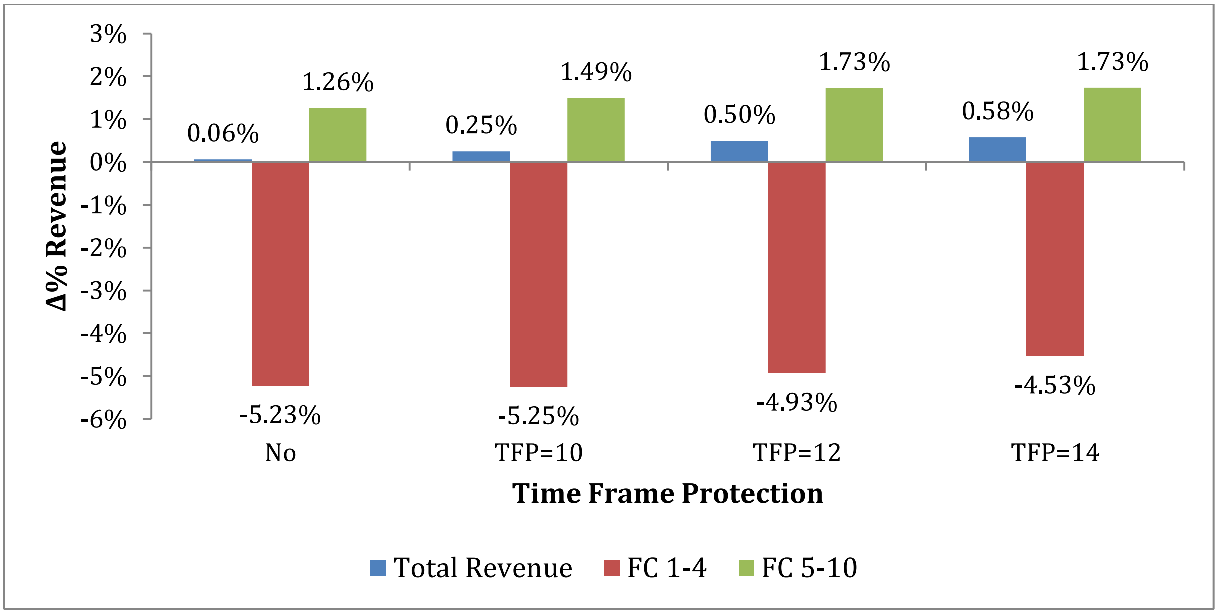

Figure 5-58: Time Frame Protection - Method 3 - Low Demand

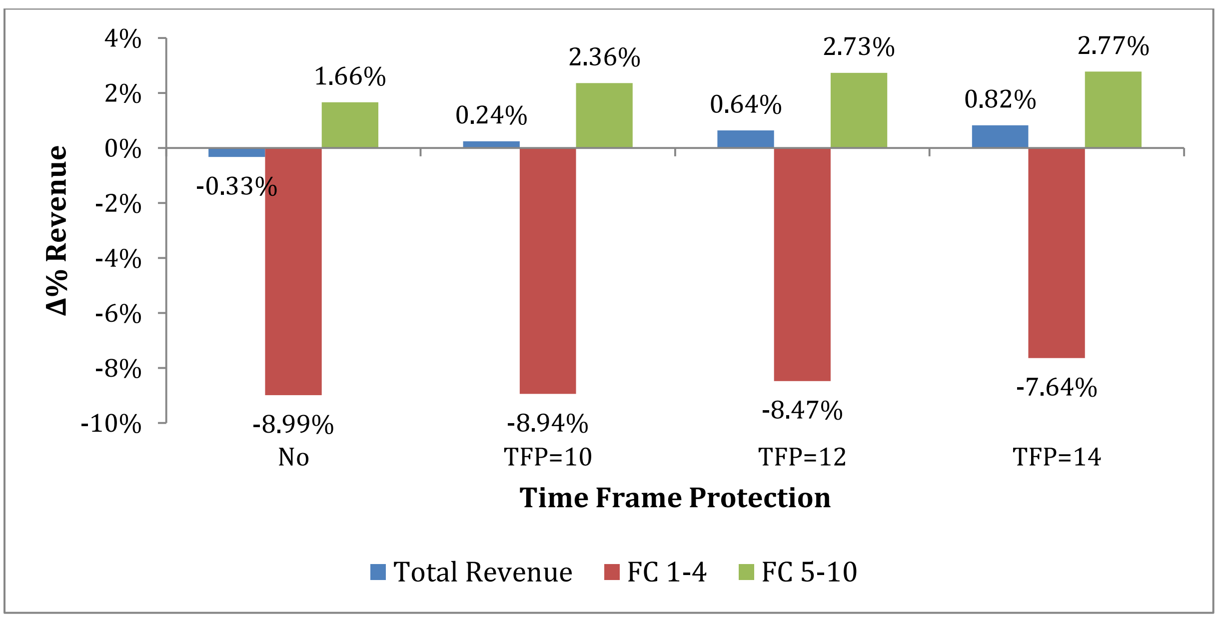

Figure 5-59: Time Frame Protection - Method 3 - Medium Demand

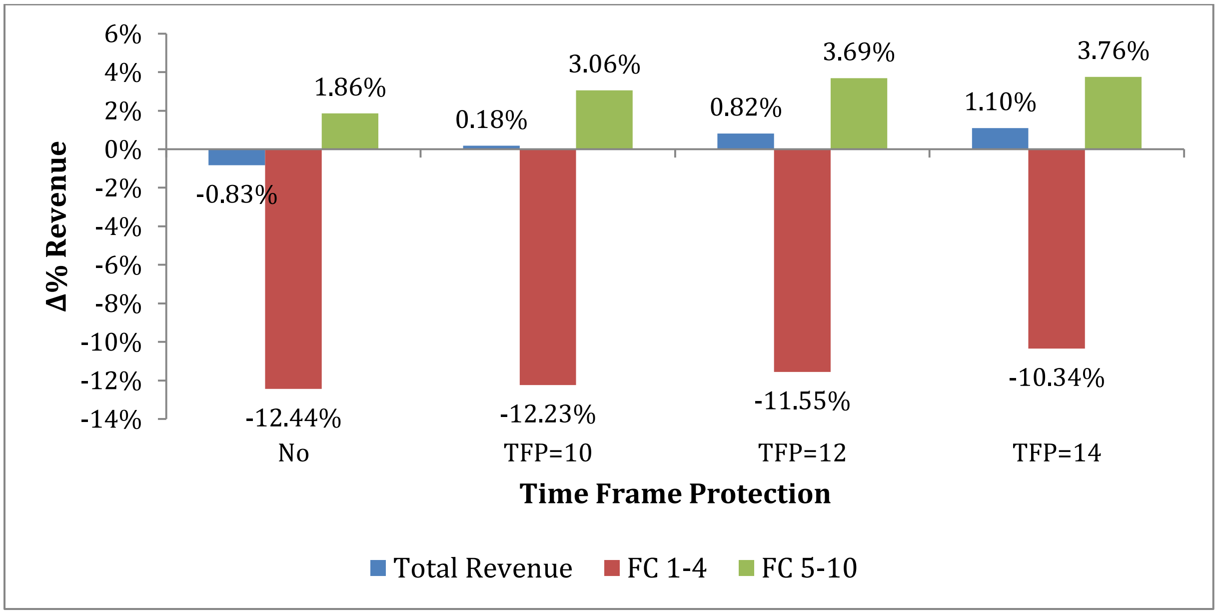

Figure 5-60: Time Frame Protection - Method 3 - High Demand

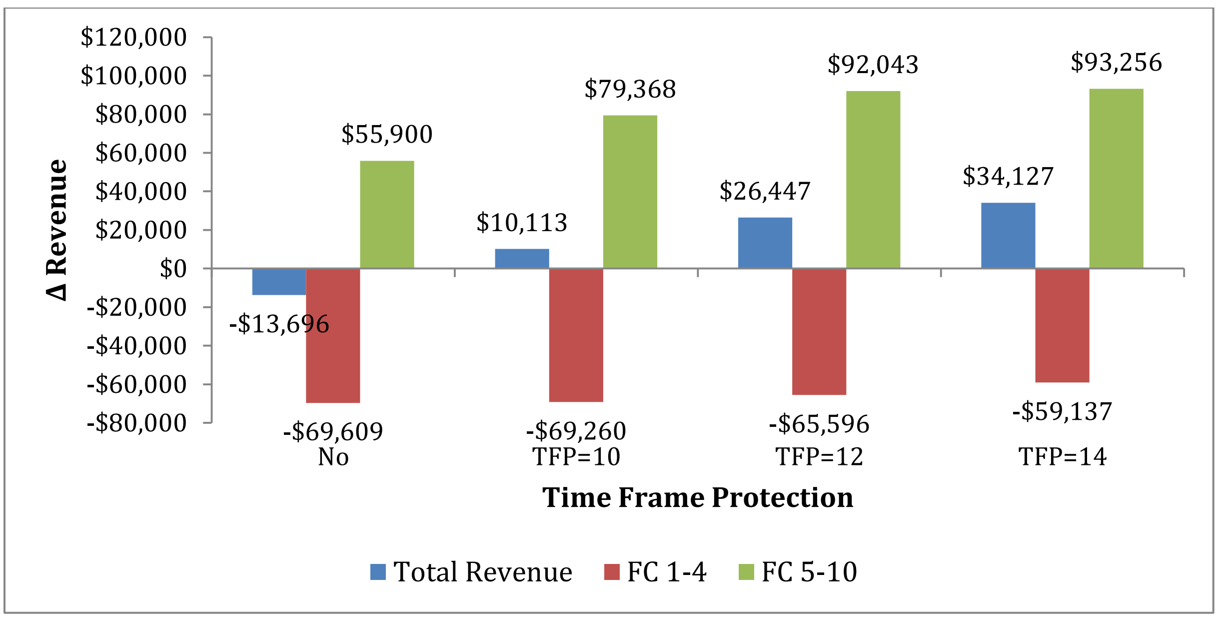

Figure 5-61: Method 3 with TFP - Absolute Revenue Gains by Fare Class Type – Medium Demand

- **Leg Load Factor Criterion Protection Mechanism**

It should be noted that Method 1 and Method 2 do not apply for this kind of protection because their corresponding heuristics set capacity sharing to begin only after economy fare class is sold out (leg load factor =100%); therefore, LLFC is only tested for Method 3 at the following levels: 70%, 80%, 90%, 95% and 99%.

Similar to the results presented for TFP, Figure 5-62, Figure 5-63 and Figure 5-64 illustrate how LLFC helps Method 3 to achieve total revenue gains; these figures also show that the highest total revenue gains are achieved with LLFC of 95%-99%.

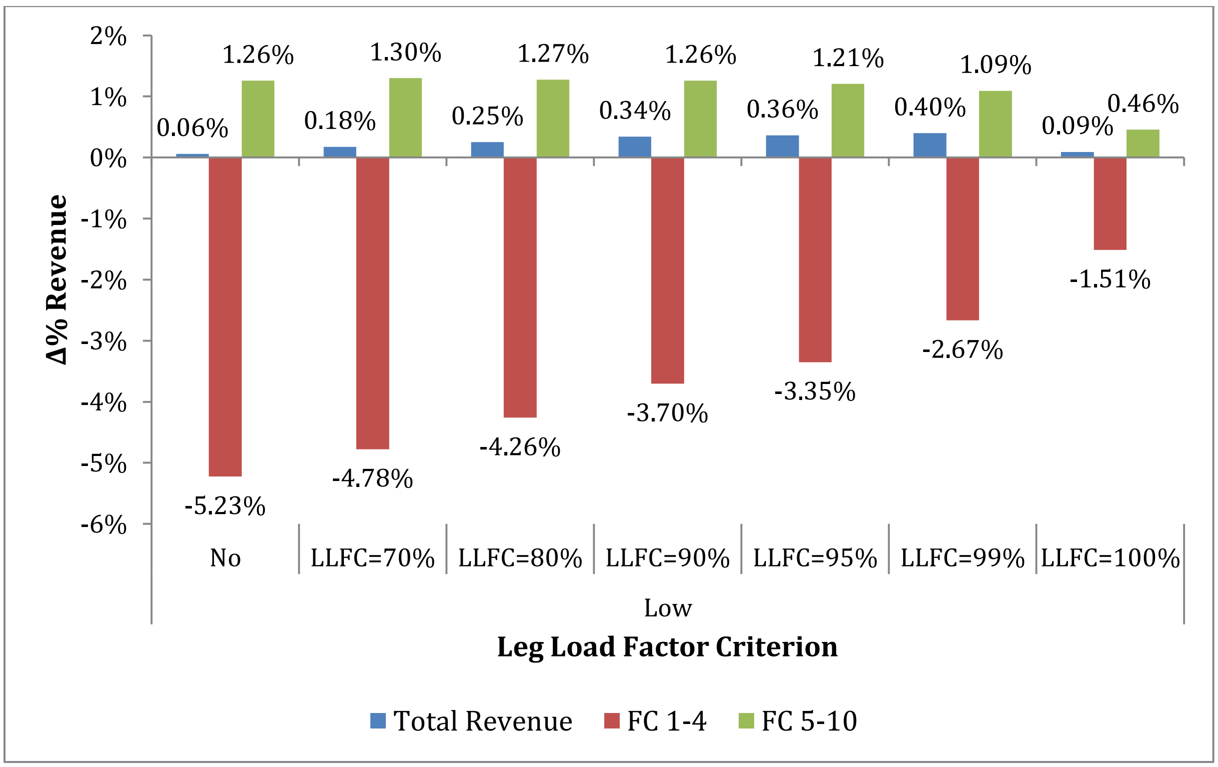

Figure 5-62: Leg Load Factor Criterion - Method 3 - Low Demand

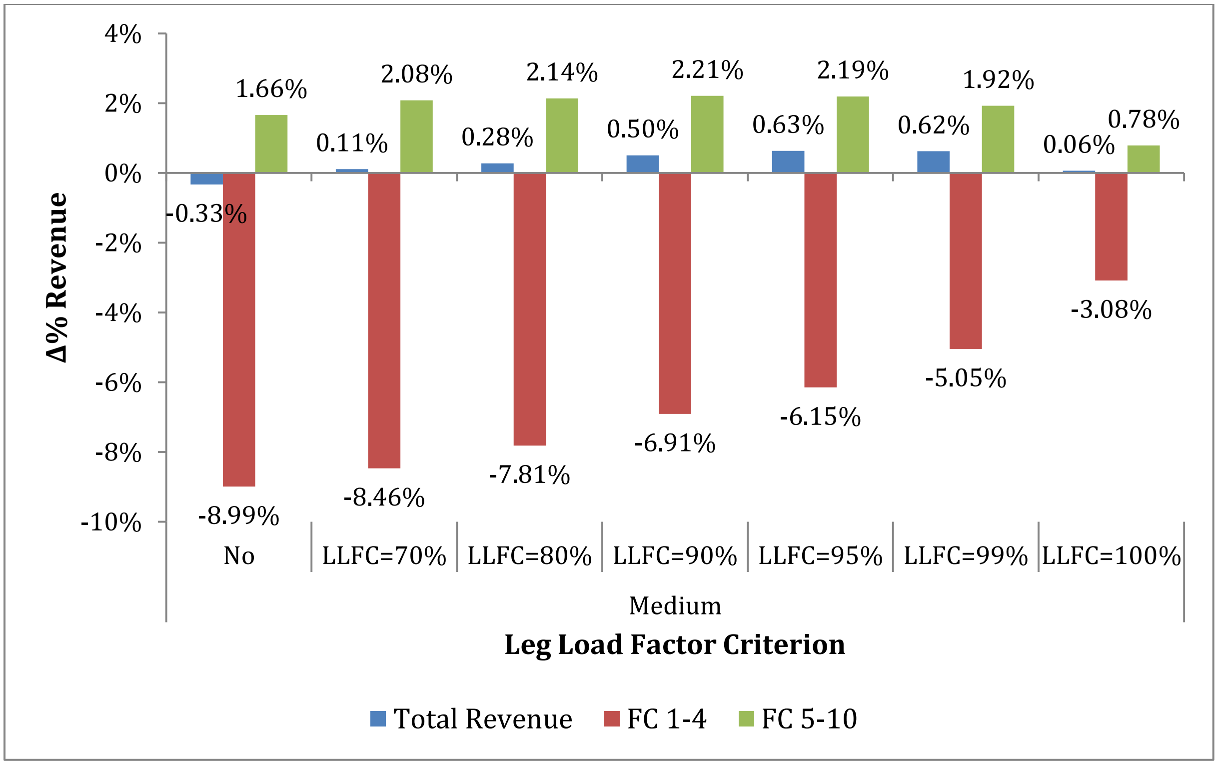

Figure 5-63: Leg Load Factor Criterion - Method 3 - Medium Demand

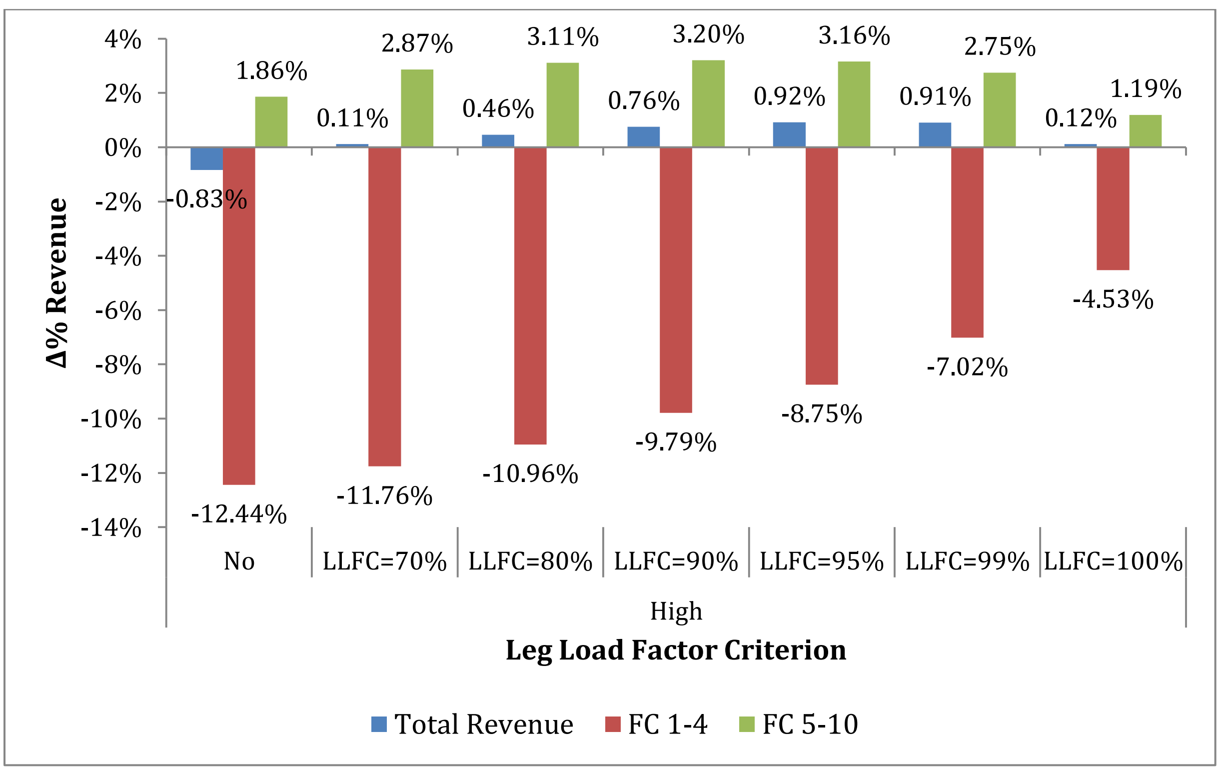

Figure 5-64: Leg Load Factor Criterion - Method 3 - High Demand

The distribution of the gains and losses between fare class groups and cabins changes in a similar way as it does when TFP is used: with high LLFC values there is less premium cabin capacity shared, so the revenue gains from passengers that book in FC 5 to 10 but are carried in premium cabin decrease together with the losses from passengers that book in premium fare classes. However, at the levels where the highest revenue gains are reported, there is a significant increase in the revenue captured from passengers flying in economy cabin with respect to the baseline (Figure 5-65).

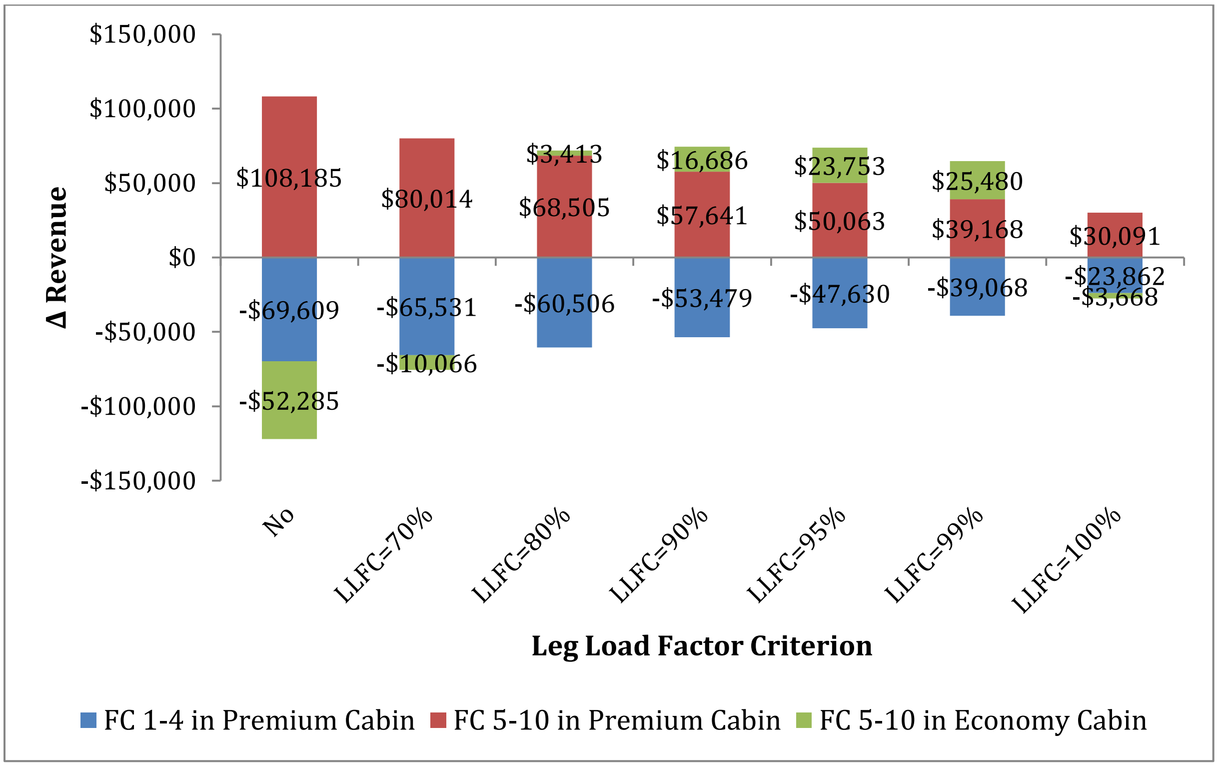

Figure 5-65: Method 3 - Absolute Revenue Variation by FC Type and Cabin Accommodation and LLFC – Medium Demand

5.3 Summary

Four premium cabin capacity sharing mechanisms have been evaluated in this chapter. The first of these, Full EMSR, is proposed for leg-control RM, and the other three mechanisms (Premium Cabin Bid Price, Joint Cabin Bid Price and Effective Bid Price) are designed for OD-control RM. All the mechanisms proposed share the same pattern: revenue losses in premium fare classes in exchange for gains in the economy fare classes. However, the mechanisms that activate the option of accommodating passengers with bookings in economy fare classes in premium cabin only after the economy cabin is sold out (i.e. Method 1 and Method 2) are the only ones that are able to generate substantial total revenue gains. In contrast, the methods that have in common that it is possible to share premium cabin capacity with economy fare classes throughout the booking period (i.e. Method 3 and Full EMSR) obtain either losses or negligible gains when capacity is shared.

In order to reduce the losses of the premium fare classes and to delay the premium cabin capacity sharing capability, two protection mechanisms were proposed: Time Frame Protection ("TFP") and Leg Load Factor Criterion ("LLFC"). While the former was tested with the four capacity sharing strategies proposed herein, the latter was only tested with the Effective Bid Price method and with Full EMSR, as it does not apply to the methods that are activated when economy cabin is sold out. With the right inputs, these protection mechanisms prove to be effective to reduce the losses in premium fare classes and to obtain gains from Method 3 and Full EMSR.

Table 5-13 summarize the ranges of proportional variations in revenue obtained with each of the OD-control premium cabin capacity sharing methods, with and without the protection mechanisms in place. Overall, the best results are obtained with the Joint Bid Price method (Method 2); in such case, using or not TFP results in very similar performance. When using TFP, similar total revenue gains are observed when using either Method 1 or Method 3; however, Method 2 is considered better because of smaller losses in premium fare classes.

	Premium Cabin BP (Method 1)	Joint Cabin BP (Method 2)	Effective BP (Method 3)
No Protection	**Total: 0.78%** FC 1-4: -7.31% FC 5-10: 2.64%	**Total: 0.86%** FC 1-4: -7.16% FC 5-10: 2.71%	**Total: -0.33%** FC1-4: -8.99% FC 5-10: 1.66%
Time Frame Protection	**Total: 0.78%** FC 1-4: -7.31% FC 5-10: 2.64% **TF=12**	**Total: 0.91%** FC 1-4: -7.19% FC 5-10: 2.77% **TF=12**	**Total: 0.82%** FC1-4: -7.64% FC 5-10: 2.77% **TF=14**
LLFC	N/A	N/A	**Total: 0.63%** FC1-4: -6.15% FC 5-10: 2.19% **LLFC=95%**

Table 5-13: Summary of proportional revenue variations for OD-control premium cabin capacity sharing methods at the Medium Demand Level

Table 5-14 presents the results for Full EMSR, showing that the implementation of this method is not beneficial without considering a protection mechanism that delays the activation of premium cabin capacity sharing. More specifically, the best performance of Full EMSR is achieved when the sharing capability in the last few time frames of the booking process.

	Full EMSR
No Protection	**Total: -0.19%** FC 1-4: -10.20% FC 5-10: 2.05%
Time Frame Protection	**Total: 0.77%** FC 1-4: -8.42% FC 5-10: 2.82% **TF=14**
LLFC	**Total: 0.59%** FC1-4:-8.64% FC 5-10: 2.66% **LLFC=95%**

Table 5-14: Summary of proportional revenue variations for Full EMSR

6. Conclusions

This chapter summarizes the motivations and context of this thesis and reviews its findings and contributions. Finally, potential directions for future research are suggested.

6.1 Dual Cabin Aircraft Capacity Management

Aircraft compartments are often divided by airlines into cabins with different service standards to satisfy the expectations of passengers with different willingness to pay and service expectations. The services provided in premium cabins typically include premium food and drinks, additional space, more comfortable seats and service oriented flight attendants. Together with the offer of a wide variety of fare products, these cabins are used by airlines to implement differential pricing, a critical concept in airline revenue management.

Once prices for these products are defined by an airline, revenue management tools are used to maximize the revenue captured from the flights in the network. This means managing the seat inventory of the aircraft such that seats are protected for travelers that arrive late in the booking process and are willing to pay higher fares. This thesis focuses on the case of dual cabin aircraft, in which the aircraft is divided into premium and economy cabins. This division of the aircraft in two sections represents a restriction to the process of revenue management, since it limits the amount of seats to be protected for the premium fare classes from the lower fare classes to the physical capacity of the premium cabins. In addition, when cabins are managed distinctly, airlines cannot use empty premium cabin seats for accommodating excess demand on the economy cabin.

The fleet decisions made by airlines impact their performance many years ahead of such decisions. Because of this, any degree of flexibility for adjusting the capacity to better match the demand provides an opportunity for airlines to maximize revenues when conditions have changed. Two options were explored in this thesis as mechanisms to provide such degree of flexibility to the airlines: first, adjusting the capacity available in each cabin by either replacing an aircraft type used by the airline for another aircraft type with a different configuration or by retrofitting the aircraft to modify its configuration; second, by allowing passengers with bookings in economy fare classes to be accommodated in the premium cabin when capacity in such cabin is available. These mechanisms are studied in this thesis and described briefly in the following sections.

6.2 Insights from Cabin Configuration Analysis

The Boeing-Swan Spill Model is frequently used by airlines as an analytical method to evaluate the impact of changing aircraft sizes on revenues. This approach is extended in this thesis to the evaluation of cabin configuration of dual cabin aircraft. Moreover, the research in this thesis was conducted by applying this adaptation of the Boeing-Swan Spill Model ("BSM") for dual cabin aircraft and comparing its results with the output of simulations performed in the Passenger Origin Destination Simulator (PODS). The analysis was performed on three different aircraft types operated by an airline, with four alternative configurations for each of these aircraft types. In addition, the simulations considered different demand levels and included scenarios where the airline was performing either leg-based or OD-control in its revenue management systems.

The results for the comparison between the BSM and PODS for aircraft type 'A' (with a baseline configuration of 150 seats in economy cabin and 20 in premium cabin), 'B' (200 in premium, 24 in economy) and 'C' are summarized in Table 6-1, Table 6-2 and Table 6-3, respectively. The initial finding was that the dual cabin BSM is able to predict reasonably well whether a change of the configuration of an aircraft will generate revenue gains or losses for an airline. As a matter of fact, only in three cases (out of seventy two scenarios evaluated) were the results of the BSM in a different direction from the PODS results; such cases are marked in dark gray in the tables below. Besides, it was also found that the BSM tends to exaggerate the magnitude of the gains and losses; this is the case for fifty nine of the scenarios evaluated (the exceptions are marked in light gray in the tables).

		Leg-RM				**OD Control**			
Demand	**Mechanism**	**126-28**	**138-24**	**162-16**	**174-12**	**126-28**	**138-24**	**162-16**	**174-12**
Low	BSM	-3.68%	-1.46%	0.70%	0.83%	-3.90%	-1.56%	0.78%	1.00%
	PODS	-2.81%	-1.38%	0.36%	0.13%	-3.32%	-1.59%	0.38%	0.15%
Medium	BSM	-5.43%	-2.34%	1.56%	2.45%	-5.05%	-2.13%	1.34%	2.01%
	PODS	-4.36%	-2.17%	1.08%	1.37%	-4.34%	-2.39%	1.01%	1.33%
High	BSM	-6.77%	-3.09%	2.44%	4.15%	-6.43%	-2.85%	2.06%	3.34%
	PODS	-4.89%	-2.59%	1.57%	2.50%	-3.90%	-1.88%	2.13%	3.02%

Table 6-1 Revenue Variation on legs operated by Aircraft Type A

		Leg-RM				OD Control			
Demand	Mechanism	168-36	184-30	216-18	232-12	168-36	184-30	216-18	232-12
Low	BSM	-2.78%	-0.97%	-0.48%	-2.75%	-3.21%	-1.20%	-0.30%	-2.49%
	PODS	-1.47%	-0.35%	-0.58%	-2.41%	-1.44%	-0.37%	-0.49%	-2.03%
Medium	BSM	-3.69%	-1.29%	-0.39%	-2.51%	-4.35%	-1.65%	-0.06%	-1.95%
	PODS	-1.87%	-0.47%	-0.43%	-1.91%	-1.74%	-0.60%	-0.29%	-1.49%
High	BSM	-4.92%	-1.89%	0.26%	-1.18%	-5.60%	-2.25%	0.54%	-0.70%
	PODS	-2.93%	-1.01%	-0.06%	-1.15%	-2.41%	-0.88%	0.06%	-0.85%

Table 6-2: Revenue Variation on legs operated by Aircraft Type B

		Leg-RM				OD Control			
Demand	Mechanism	218-42	234-36	266-24	282-18	218-42	234-36	266-24	282-18
Low	BSM	2.36%	1.51%	-2.19%	-4.81%	1.35%	1.06%	-1.81%	-4.12%
	PODS	0.88%	0.66%	-1.32%	-3.33%	0.76%	0.64%	-1.21%	-3.08%
Medium	BSM	1.94%	1.25%	-1.92%	-4.20%	0.79%	0.77%	-1.51%	-3.56%
	PODS	0.63%	0.56%	-1.20%	-2.91%	0.58%	0.55%	-1.13%	-2.82%
High	BSM	1.49%	0.97%	-1.60%	-3.52%	-0.90%	-0.13%	-0.54%	-1.65%
	PODS	0.41%	0.52%	-1.09%	-2.57%	0.38%	0.46%	-1.10%	-2.46%

Table 6-3: Revenue Variation on the legs operated by Aircraft Type C

The main reason identified for the differences between the BSM and the results found in PODS is that the BSM assumes that the demand for travel in one cabin is independent from the demand for travel in the other cabin. By contrast, the simulations in PODS represent a significant level of interaction between the bookings in premium fare classes and economy fare classes. This has an effect on both the estimated load factors for each cabin and the additional passenger value of the passengers spilled or accommodated when changing the configuration of the aircraft.

6.3 Insights from Premium Cabin Capacity Sharing

In order to reduce the impact on revenues of the constraint imposed by the division of the aircraft into sub-cabins, premium cabin capacity sharing schemes were proposed for taking advantage of available premium cabin capacity to accommodate additional passengers with bookings in economy fare classes. The objective of this mechanism is to increase the total revenue captured by an airline by offering premium cabin seats expected to be empty to passengers booking in economy fare classes (fare classes 5 to 10). The expectation of the premium cabin capacity sharing schemes is increasing the revenue generated from the premium cabin by adding passengers with bookings in economy fare classes to the passengers that are expected to buy premium fare classes (fare classes 1 to 4).

"Full EMSR" is proposed as the premium cabin capacity sharing mechanism for leg-RM. This methodology considers applying the leg-based revenue management heuristic EMSRb to the entire aircraft capacity and modifying the booking limits for premium fare classes based on the capacities of each cabin. As summarized in Table 6-4, the results of the simulations of Full EMSR in PODS show that this mechanism does not behave as expected; premium cabin capacity sharing using Full EMSR results either in a decrease or just a slight increase in total revenue (premium + economy).

Concept	Expected	Simulations
Total Revenue (premium + economy)	Increases	Increases slightly/Decreases
FC 1-4 Revenue	Unchanged	Decreases
FC 5-10 Revenue	Increases	Increases
Premium Cabin Revenue	Increases	Increases
Economy Cabin Revenue	Unchanged	Decreases

Table 6-4: Comparison of Expected and Observed Results – Full EMSR

It was found that the performance of Full EMSR is not beneficial because premium cabin capacity is shared with economy fare classes from the beginning of the booking process. This results in more bookings in economy fare classes and less in premium fare classes, affecting the forecasts used for determining the protection levels for the highest fare classes (similar to the spiral down concept).

Two mechanisms were proposed to delay the activation of premium cabin capacity sharing. The first of these is named Time Frame Protection, as it allows capacity to be shared only after a predetermined time frame in the booking process. The second mechanism is based on a Leg Load Factor Criterion, where premium cabin capacity is shared only after the economy cabin booked load factor is higher than a predetermined level. As presented in Table 6-5, total revenues using Full EMSR are increased when these protection mechanisms are in place; the 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 (TF=14).

	Full EMSRb
No Protection	**Total: (0.03%- -0.66%)** FC 1-4: (-5.81% - -14.59%) FC 5-10: (1.31%- 2.54%)
Time Frame Protection	**Total: (0.51 - 1.06%)** FC 1-4: (-4.96%--11.7%) FC 5-10: (1.72%- 3.98%) **TF=14**
Leg Load Factor Criterion	**Total: (0.46%-0.78%)** FC1-4:(-4.95%--11.94%) FC 5-10: (1.65%-3.70%) **LLFC=95%**

Table 6-5: Summary of proportional revenue variations for Full EMSR

In the case of OD-control RM, three mechanisms were proposed: the first of these is named "Premium Cabin BP" or simply "Method 1" and allows premium cabin capacity to be shared only after economy cabin is sold out. The criterion for accepting bookings from economy fare classes in premium cabin is by comparing the fare of the requested economy fare class with the sum of the premium cabin bid price of the legs traversed by the ODF. If the fare is higher than the applicable premium cabin BP, then the booking is accepted.

The second OD-control premium cabin capacity sharing mechanism is named "Joint Cabin BP" or "Method 2". In this case a joint cabin bid price is calculated for the complete aircraft capacity (in addition to the bid price applicable for each cabin), ignoring the division between premium and economy fare cabins. After economy cabin capacity is sold out, the criterion used for accepting bookings from economy fare classes in premium cabin is by comparing the fare of the requested economy fare class with the sum of the joint cabin bid price of the legs traversed by the ODF. If the fare is higher than the applicable joint cabin BP, then the booking is accepted.

Table 6-6 compares the results of the simulations of Method 1 and Method 2 in PODS with the expected behavior of the premium cabin capacity mechanisms. In this case, gains in total revenue are effectively observed; however, these gains are not generated from the expected sources. For example, instead of remaining unchanged, the revenue from premium fare classes decrease as a consequence of sharing premium cabin capacity (typically a very negative outcome for an airline);

on the other hand, the revenue captured from the economy cabin increases although it was not an expected outcome of the implementation of these methods.

Concept	Expected	Simulations
Total Revenue (premium + economy)	Increases	Increases
FC 1-4 Revenue	Unchanged	Decreases
FC 5-10 Revenue	Increases	Increases
Premium Cabin Revenue	Increases	Mixed – Sometimes increases, sometimes decreases
Economy Cabin Revenue	Unchanged	Increases

Table 6-6: Comparison of Expected and Observed Results – Method 1 and Method 2

A third method, denominated as "Effective BP" or "Method 3" allows for premium cabin capacity sharing throughout the booking process regardless of the load factor in any of the cabins. In this case a joint bid price for all the aircraft is calculated (in addition to the bid prices for each of the cabins) and is considered for accepting bookings in both cabins. More specifically, the bid price that is effectively used for accepting or rejecting premium fare classes' bookings is the maximum of the premium cabin bid price and the joint cabin bid price. By contrast, the effective bid price used for opening or closing economy fare classes is the minimum of the economy cabin bid price and the joint cabin bid price.

Table 6-7 summarizes the performance of the Effective BP method, with the main finding that total revenue decreases or increases slightly. As observed with Full EMSR in the case of Leg-RM, this mechanism negatively affects the revenue from premium fare classes as well as the revenue captured in the economy cabin.

Concept	Expected	Simulations
Total Revenue (premium + economy)	Increases	Increases slightly/Decreases
FC 1-4 Revenue	Unchanged	Decreases
FC 5-10 Revenue	Increases	Increases
Premium Cabin Revenue	Increases	Increases
Economy Cabin Revenue	Unchanged	Decreases

Table 6-7: Comparison of Expected and Observed Results – Method 3

Considering the findings with the protection mechanisms presented above for Full EMSR, the same approach was used for the OD-control capacity sharing methods described. It should be noted that the LLFC method does not apply to Methods 1 and 2 since capacity is shared in those methods only after the economy cabin is sold out.

The results obtained from the PODS tests for each of the OD-control premium capacity sharing methods described above are summarized in Table 6-8. It is shown that the Time Frame Protection mechanism does not have a significant effect on the results of Methods 1 and 2, since these methods are already sharing premium cabin capacity late in the booking process by imposing the condition of having economy cabin sold out before activating capacity sharing. On the other hand, the protection mechanisms are useful to improve the performance of the Effective BP method because of delaying capacity sharing until the last time frames of the booking process. Overall, the Joint Cabin BP method seems to provide the highest total revenue gains between the methods explored in this thesis. The performance of the other methods is similar when protection mechanisms are also considered, but the revenue losses in premium fare classes are still lower in Method 2, making it a more beneficial mechanism.

	Premium Cabin BP (Method 1)	Joint Cabin BP (Method 2)	Effective BP (Method 3)
No Protection	**Total: (0.54%- 1.08%)** FC 1-4:(-4.15% - -9.91%) FC 5-10: (1.6%- 3.63%)	**Total: (0.58%-1.16%)** FC 1-4: (-4.11%--9.76%) FC 5-10: (1.65%- 3.7%)	**Total:(-.06%- -.83%)** FC1-4:(-5.23%--12.44%) FC 5-10: (1.26%- 1.86%)
Time Frame Protection	**Total: (0.54 - 1.08%)** FC 1-4: (-4.15%--9.93%) FC 5-10: (1.6%- 3.64%) **TF=127days**	**Total: (0.58 - 1.16%)** FC 1-4: (-4.11%--9.76%) FC 5-10: (1.65%- 3.7%) **TF=12 – TF=10**	**Total: (0.58%- 1.1%)** FC1-4:(-4.53%--10.34%) FC 5-10: (1.73%- 3.76%) **TF=14**
Leg Load Factor Criterion	N/A	N/A	**Total: (0.4%-0.92%)** FC1-4:(-2.67%--8.75%) FC 5-10: (1.09%-3.16%) **LLFC=95% - LLFC=99%**

Table 6-8: Summary of proportional revenue variations for OD-control Premium Cabin Capacity Sharing Mechanisms

6.4 Suggestions for Future Research

Many possibilities exist for continuing research on capacity management for multi-cabin aircraft. The first of these is related with the extension of this research to aircraft with more than two cabins. As part of the differential pricing strategy of some airlines, products beyond the typical economy and premium cabins are being offered nowadays; some examples of these cabins are Premium Economy (with service standards between Economy and Premium), First Class (with higher service standards than in Premium) and Suites (with better service than First Class). As discussed before, the interaction between the passengers of both cabins observed in PODS but ignored by the BSM explain many of the differences between the results obtained when using those methods to evaluate a potential aircraft configuration change. Furthermore, the relationship between the cabins is also observed when premium cabin capacity is shared, as the revenue obtained from premium fare classes always decrease as a result of the increase on the forecasts for bookings in economy fare classes. With additional cabins in the aircraft operated by the airline these interactions between passengers may become more difficult to understand and to analyze using an analytical model. Therefore, a simulation tool such as PODS can provide powerful insights that allow understanding the cabin configuration analysis in aircraft with more than two cabins. Similarly, the case of capacity sharing the mechanisms presented in this document may be studied in the case of aircraft with three or more cabins.

In this thesis 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.

Another potentially interesting field for research consists on considering simultaneously the cabin configuration analysis and the premium cabin capacity mechanisms proposed in this thesis. Based on the results presented in Chapter 4, in certain scenarios airlines would not benefit from adding premium cabin seats to an aircraft type because of the risk of rejecting many passengers that would have booked in economy fare classes. As a matter of fact, an airline can observe that adding premium cabin capacity might generate gains on some legs operated by an aircraft type but losses on the rest of the other legs. Nevertheless, premium cabin capacity sharing provides an additional degree of flexibility to the airline as it can offer premium cabin seats to passengers flying in economy fare classes on legs with low demand for premium cabin seats, and dedicate its premium cabin capacity to premium fare classes on flights with high demand for premium cabin seats. Hence, premium cabin capacity sharing could potentially lead to better results for the configurations that increase the number of seats in premium cabin for every aircraft type.

The results of Chapter 5 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, it should be noted that as part of the process of setting the parameters for the simulations in PODS, a set of disutilities are assigned to each passenger based on predetermined mean values that lead to a well calibrated baseline network. However, the variation of the values of these parameters could potentially impact the results of the tests presented in this thesis. For example, if the dis-utility of the average passenger for flying in premium cabin is higher than the value considered for the runs included herein, it would be expected that more people would be willing to fly in premium cabin. Moreover, those values could modify the degree of dependence of the demand for each of the cabins; considering how important was the interaction between passengers of both cabins in the results presented in this thesis, specific research could be dedicated to understand how strong (or weak) are the disutilities of traveling in economy cabin for some passengers and in some markets.

Bibliography

[1] Abramovich, M. (2013). *Impacts of Revenue Management on Estimates of Spilled Passenger Demand*. Cambridge, MA. Master's Thesis, Massachusetts Institute of Technology

[2] 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.

[3] Boeing Commercial Airplane Company, (1978) *Load Factor Analysis: The Relationship Between Flight Load and Passenger Turnaway*, Working Paper, Seattle, Washington

[4] Belobaba, P. (1987) *Air Travel Demand and Airline Seat Inventory Management.* PhD thesis, Massachusetts Institute of Technology.

[5] Belobaba, P. (1992) *Optimal vs. heuristic methods for nested seat allocation*. In AGIFORS Reservations Control Study Group Meeting.

[6] Belobaba P. (2016) *Optimization Models in RM systems: Optimality versus revenue gains*. Journal of Revenue and Pricing Management. Vol. 00, 0, 1-7 (forthcoming)

[7] Belobaba, P. P., Farkas, A. (1999). *Yield Management Impacts on Airline Spill Estimation*. Transportation Science, 33(2), 217-232.

[8] Belobaba, P., Odoni, A., Barnhart, C. (2015), *The Global Airline Industry*, Wiley, 2015

[9] Bratu, S. (1998). *Network Value Concept in Airline Revenue Management.* Cambridge, MA: Master's Thesis, Massachusetts Institute of Technology.

[10] Farkas, A. (1996), *The Influence of Network Effects and Yield Management on Airline Fleet Assignment Decisions*, PhD. dissertation, Flight Transportation Laboratory Report R96-1, MIT, Cambridge, Massachusetts

[11] Fry, D.G. (2015), *Demand Driven Dispatch and Revenue Management*. Cambridge, MA. Master's Thesis, Massachusetts Institute of Technology

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

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

[14] Li, M.Z.F., Oum, T. H., (2000), *Airline spill analysis - beyond the normal demand*. European Journal of Operational Research 125 p. 205-215

[15] Nita, M., Scholz, D. (2012), *Business opportunities in aircraft cabin conversion and refurbishing*, Journal of Aerospace Operations 1 (2012) 129–153

[16] Subramanian, R., Scheff, R., Quillinan, J, Wiper, D.S., Marsten, R. (1994), *Coldstart: Fleet Assignment at Delta Air Lines*. Interfaces (24:1) pp. 104-120.

[17] Swan, W.M. (1994) *Using the Spill Model*, Working Paper. Boeing Marketing

[18] Swan, W.M. (2001) *Spill Modeling for Airlines*. Boeing Marketing

[19] Swan, W.M. (2002) *Airline demand distributions: passenger revenue management and spill*. Transportation Research Part E, V 38, p. 253-263

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

[21] Walczak, D., (2010) Capacity sharing. In *AGIFORS*. 2010 PROS Holdings, Inc.