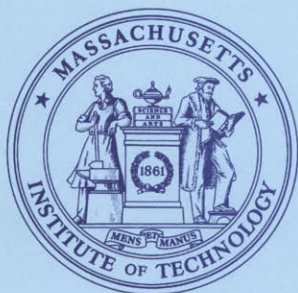FLIGHT TRANSPORTATION LABORATORY
REPORT R93-4

AIRLINE PRICING AND FARE PRODUCT
DIFFERENTIATION

THEODORE CHARLES BOTIMER

DEPARTMENT OF

AERONAUTICS AND **ASTRONAUTICS**

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

CAMBRIDGE, MASSACHUSETTS 02139-4307

AIRLINE PRICING AND FARE PRODUCT DIFFERENTIATION
by
THEODORE CHARLES BOTIMER

Submitted to the Department of Civil and Environmental Engineering
on October 20, 1993 in partial fulfillment
of the requirements for the Degree of
Doctor of Philosophy in Transportation Economics

ABSTRACT

Airlines offer combinations of price level and purchase restrictions, or fare products, designed to best maximize revenues on their flights. This dissertation provides the first comprehensive examination of the differentiated fare product structure on the market today in terms of passenger demand, airline revenue, and societal welfare. The role of pricing in the airline revenue management function is established first. The types of price discrimination currently practiced by airlines, as defined in the economics literature, are then identified. Although the terms airline product differentiation and price discrimination have previously been used interchangeably, the two practices are shown to differ and exist simultaneously in the current industry environment.

Next, airline revenue management techniques and, in particular, fare product differentiation are examined from the standpoint of economic efficiency. This dissertation concludes that both efficiency in exchange and Pareto optimality are unattainable under the current structure of airline fare product differentiation as a result of the costs incurred by passengers due to applied purchase restrictions. It is found, however, that a differentiated fare product structure with a wide range of price levels coupled with effective revenue management techniques can provide airline seats to those consumers who value them most when demand exceeds supply. Efficiency in allocation can thus be achieved in the current industry environment.

Virtually every existing yield management seat allocation model assumes that consumers view differentiated airline fare products as separate products with

2

uncorrelated demands that compete for space on a fixed capacity aircraft. Such formulations ignore the dependence of the demand for a given fare product on the price levels and characteristics of the other available (competing) fare products. In this dissertation, a model of product differentiation that considers the interrelationships of the available airline fare products as well as the cost incurred by consumers of accepting more restricted (and less flexible) products is presented. This generalized cost model of airline fare product differentiation explicitly incorporates the techniques of fare product differentiation and price discrimination currently used by airlines.

The generalized cost model is extended to incorporate the "buy down" or diversion of passengers to lower-priced fare products as a result of their ability to meet the additional purchase restrictions imposed by airlines. Moreover, diverting passengers may be induced to "sell up" to higher-priced fare products when booking limits are applied to the lower-priced products. The generalized cost model contributes the first behavioral motivation of both passenger diversion and sell up. The dissertation demonstrates the use of booking limits as devices to control and limit the revenue dilution effects of passenger diversion.

The effects of pricing and other fare product design decisions are quantified for any set of OD market conditions using the generalized cost model. The model provides insight into the underlying effects of the tradeoffs made by airlines when making pricing and marketing planning decisions. In summary, this research provides the first cohesive look at the relationships between price level, purchase restrictions, demand, and revenue in the context of airline product differentiation and yield management.

Thesis supervisor: Peter P. Belobaba
Title: Assistant Professor of Aeronautics and Astronautics

Table of Contents

List of Figures

List of Tables

Chapter One

Airline Fare Product Differentiation and Revenue Management

The travel options that are available to consumers contemplating the purchase of airline services have grown extensively following the Airline Deregulation Act of 1978 which gave U.S. carriers complete control over the pricing and supply of domestic air transportation services sold. Similar to the U.S. domestic trend, the number of air travel options available to consumers around the world has also grown. The trend in international markets has been less dramatic, in general, since fare changes require bilateral government approval.

The large number of carrier choices in each market complicates matters further when combined with the many service choices. Increases in carrier choices have been fueled by the development of large hub and spoke route structures which enable airlines to offer service in more numerous origin to destination (OD) markets. In addition to the plethora of published fares that exist in the computer reservations systems (CRS) of the airlines for immediate sale to the public, all airlines are willing to negotiate different price levels with important consumers or volume buyers such as group tour operators, increasing the number of air travel options further.

The increases in air travel options have not come without cost since the vast array of choices available to the traveling public and the ability of airlines to change the structure of the services that they offer has lead to confusion among many consumers. This confusion stems from the large number of options available and the difficulties that most passengers have in obtaining up-to-date travel information, even with the assistance of travel agents. In the eyes of many, the primary result of the service provision flexibility given to airlines has been an overcomplication of available air travel service offerings. In fact, the pricing structure that exists in the airline industry today has routinely been characterized as unfair, complicated, and incomprehensible.

The U.S. domestic industry provides an illustrative example of how complicated the air travel service structure has become. Pricing by United States domestic airlines has changed dramatically as a result of the Airline Deregulation Act of

1978. Within the coach class cabin, airline product choices have increased from a single type of service to the multiplicity of service options that are available on the market today. The complexity has increased rapidly over the fifteen years of deregulation, considering the relative stability and simplicity of airline fare setting policy during regulation.

The services purchased by the public are for travel in a specific origin to destination or OD market. The travel services available can be for one way or round trip travel within the OD market. While many consumers purchase air travel services through special negotiated deals with the airlines, the majority of consumers purchase the bundles of price level and purchase restrictions, or fare products, from the published fares listed on the CRS systems. The published fares are, in general, for scheduled airline services since the impact of charter flights in the U.S. market has declined to relative insignificance. The focus of this dissertation is the effect on air travel demand and revenue of the airline fare product structure and, in specific, the fare products offered for sale to consumers contained within the computer reservations systems. The pricing and restriction setting policies engaged in by airlines are explained and suggestions for their improvement are made.

With the new differentiated structure of airline fare products and the fixed capacity that characterizes all airline services, techniques of capacity control have been introduced to best utilize the scarce resource of available seats. The dissertation shows that the current structure of service offerings can be justified in the presence of correct capacity control. Moreover, the structure may even benefit the majority of consumers. In any case, the changes in the nature of airline services require a new way of looking at airline pricing in the face of product differentiation and yield management. From the perspective of the airline, analysis of the fare product structure requires a method of measuring the effects of the price levels and restriction setting policies on revenues and demand. In this vein, airlines are interested in knowing how to best set price levels to maximize revenue. This dissertation focuses on the characterization of pricing differentiated airline fare products in the current industry environment.

This chapter examines the evolution of airline service offerings from regulation in the U.S. to the current deregulated environment. Product differentiation as practiced by airlines is described. Next, the airline (managerial) functions that

attempt to manage demand are described. An in depth look at how pricing is done at airlines is then undertaken. Finally, an outline of the dissertation is presented to provide a foundation for the further exploration of airline pricing and product differentiation.

1.1 Differentiated Airline Fare Products

During the greater part of U.S. airline regulation, the vast majority of traffic was carried at a single fare set by the governing board of commercial aviation, the Civil Aeronautics Board (CAB). A single type of reservation was required (and a single fare paid) for travel in the coach class cabin. In essence, the airline provided a service guarantee to the consumer. Fares were set according to a mileage-based formula that was developed by the CAB. The formula sought to provide airlines with a reasonable rate of return on investment for providing air transportation service in the markets designated by the CAB. Note that the term "market", in this case and throughout the dissertation, refers to a single origin to destination (OD) market. The unit of sale in the airline industry is the individual seat in an origin to destination market. All transactions made by consumers and airlines happen at the OD market level making it the relevant reference. Prior to 1979, exclusively carriers licensed by the CAB provided service in interstate U.S. air transportation markets thus ensuring conformity to the single fare pricing structure.

The CAB fare formula guaranteed a return of over 10% on investment for all carriers, given standard industry load factors and flight operations. The fare levels received by the airlines were increased annually to cover costs. The dramatic improvement in aircraft efficiency that occurred during the course of regulation was not entirely transferred to the passenger population in the form of reduced fare levels but rather through greater service levels. Some might argue the efficiency gains were translated into airline employee benefits. Since airlines were not allowed to compete on price and were virtually guaranteed a return on investment, non-price competition resulted. Airlines offered increased service levels to passengers. In-flight amenities, frequency of service, and other passenger inducements increased under regulation while unit costs and fare levels did not.

There was incentive for CAB licensed carriers to increase capacity since there was little risk of not recovering costs and passengers were sensitive to frequency of service. The simplest solution to the problem of not rejecting demand was to increase aircraft size and frequency of flights. The excess capacity resulting from this form of non-price competition presented a challenge to the industry -- finding passengers to purchase the unused seats. Under regulation, however, the methods available to carriers to stimulate passenger demand were quite limited, although some attempts were made.

1.1.1 Historical Perspective : Surplus Seat Sale

Discount fares were offered to people in an attempt to fill empty or "surplus" seats with passengers willing to pay significantly less than the unrestricted fare. The targeted passenger populations of the lower-priced fare products were not willing to pay the full fare but would choose to purchase air travel, at the right price. The airlines attempted to fill their excess capacity by developing products that appealed to the price-sensitive travelers while not decreasing revenues received from those willing to pay the unrestricted fare.

To stimulate demand among the non-traveling population, the price of airline travel needed to be reduced. Early attempts at stimulating demand on low load factor flights in regulation included "night coach" fares which were offered on late night flights operated primarily for repositioning aircraft due to schedule inefficiencies. First introduced in the fifties, the night coach fares sought to recover at least some revenue from the repositioning flights to contribute to fixed costs since little full-fare demand existed late at night. To stimulate demand further, student fares with 50% off on a stand-by basis or 20% off on a reservation basis were introduced in the mid-sixties. Fare discounting sought to stimulate incremental demand, increase load factors, and fill excess capacity.

In 1977, near the end of U.S. airline regulation, more wide-ranging fares aimed at stimulating leisure demand were introduced to the market. The "Supersaver" offered by American Airlines was aimed at stimulating demand in the price-sensitive segments of the markets. The Supersaver fares provided a 45% discount but required a 30 day advance purchase, round trip travel and an minimum stay of seven days. The Supersaver fares were similar to the "Peanuts" discount fares, originally offered by the Texas intrastate carrier Texas Air

International, provided a 50% discount off of the unrestricted fare on a flight-specific basis in five low demand markets.

Travel and purchase restrictions were attached to the air travel services to prevent full-fare passengers from purchasing the discounted services. The discounted fare products with attached purchase restrictions enabled the airlines to stimulate traffic on their routes from passengers less sensitive to service convenience and travel flexibility without diluting the revenues earned by the airline from the more price-insensitive members of the consumer population. The intention of the fare discounting practiced during regulation was to sell off the surplus seats to make contributions to fixed costs.

Drawbacks of Discounts

The discount fare offerings were not, however, without their drawbacks. The student stand-by fares resulted in operational problems for the airlines when the students learned the system and flooded certain flights with bogus bookings to create stand-by space for themselves. This resulted in lost revenue for the airlines and decreased level of service for the full-fare passengers (Kahn, 1970).

Another problem facing airline managers was the coincidence of preferred travel times for the full-fare and discount travel markets. Many flight departures popular with full-fare passengers were also quite popular with the discount travelers. Because of capacity constraints on the peak utilization flights, the airlines were required to impose limits on the number of discount fare products sold. For example, Supersaver sales were limited to 35% of the aircraft capacity, an early application of booking limits to prevent revenue dilution. In fact, any discount passenger accepted on a flight in lieu of a full-fare passenger on the same itinerary would have a negative revenue impact for the airline. Thus, stimulation of traffic without displacement of full-fare passengers was the true goal of the airline industry.

1.1.2 Market Segmentation Through Fare Product Differentiation

The surplus seat concept has been modified somewhat during deregulation. Instead of viewing the offer of discount seats to consumers as selling off the excess capacity, airlines now view the different combinations of price levels and

restriction bundles, or fare products, as separate travel options to be marketed to different segments of the population. Airlines attempt to segment the market by offering fare products designed to identify the population by their willingness to pay for air travel.

The segmentation centers on the belief that passengers who place a higher value on air travel tend to value flexibility and guaranteed availability more than price level. Passengers placing lower values on air travel are assumed more sensitive to price level and less so to flexibility and availability. An entire spectrum of passenger types ranging from the least price-sensitive/most service-sensitive passengers to the most price-sensitive/least service-sensitive exists. To approximate the desired segmentation of demand, airlines have imposed purchase restrictions and other devices. The purchase restrictions are expected to act as "fences" to prevent passengers with higher values of willingness to pay for air travel from purchasing discount fare products. Common industry fare restrictions include :

1) advance purchase requirements
2) required Saturday night stay-over
3) blackout periods
4) peak vs. off-peak travel requirements
5) weekday vs. weekend travel requirements
6) flight validity restrictions (good for travel between...)
7) ticketing purchase restrictions (purchase tickets by...)
8) required round trip travel

The above devices have been found effective by airlines in preventing passengers with high values of willingness to pay from purchasing lower-priced fare products using traveler sensitivity to season, day of week, time of day, and trip purpose. These flexibility-based "fences" attempt to exploit the relative sensitivity of passengers to convenience and convertibility of airline travel services.

The different fare products are developed to appeal to a certain type of air traveler. For instance, the full-fare unrestricted fare products are aimed at appealing to customers who value the flexibility of an itinerary and are willing to pay a higher fare to insure ticket changeability. The deepest discount fares, on

the other hand, are targeted to consumers who are willing to plan their trip far in advance, make an extended stay in a location, and sacrifice departure time to secure the lowest available fare. The discount fare products offer reduced price levels in exchange for the acceptance of flight restrictions. The lower the price level, the more severe the restrictions.

In the ideal, airlines hope to target the passengers falling into the same range of willingness to pay and entice them to buy air travel at the highest possible fare level. In the presence of "perfect fences", the value differential between the higher-priced and lower-priced fare products is assumed to be equal to the value that the passenger places on the higher-priced fare product. In other words, the lower-priced fare products are intended to be effectively useless to the perfectly fenced passenger. If not, the passenger might have incentive to "buy down" and purchase the lower-priced fare product.

Restrictions : Purely Bad or Not?

Under fare product differentiation, any two passengers seated next to each other on an aircraft are likely to have purchased completely different fare products with different price levels and purchase restrictions. At first glance, the restrictions placed on airline fare products targeted to portions of the population with low values of willingness to pay are unambiguously negative. A closer examination will reveal, however, that this is not exclusively the case. Even though all discount fare products are less valuable than the unrestricted (in all attributes except price level), some of the discount fare products offered by airlines do provide some degree of added benefit to consumers.

The advance purchase and stay requirements can be viewed as devices used purely to segment the population by their willingness to pay with no outward benefits provided to consumers. The airlines are, in fact, making their fare products less valuable to consumers by artificially reducing the travel flexibility of all passengers who purchase the more restricted products. In the most favorable scenario, the advance purchase and minimum stay requirements can be seen as neutral to the value placed on travel by a consumer. The restrictions limiting the flexibility of purchase and travel decrease product quality.

The absence of refundability penalties, on the other hand, can be viewed as a bonus to consumers since most sellers of goods and services do not allow full refundability for the cancellation of a reservation. Airlines allow full-fare and most non-excursion (*e.g.* not requiring a minimum stay) fare product passengers to arrive after the aircraft has departed and provide them with a full refund regardless of whether the seat that they reserved resulted in lost revenue or not. A similar deal is not found for theatrical and sporting events which offer no refunds whatsoever or for hotels which require, in general, at least several hours cancellation notice to receive a refund. Complete refundability is quite a generous offer for a seller with such a highly perishable product. Hence, there is value added to every refundable ticket issued by an airline demonstrating true product differentiation rather than just product degradation.

Finally, an attribute of the differentiated airline fare product that can be either beneficial or detrimental to passengers is availability. Availability limits are beneficial to the higher-fare passengers to the detriment of discount travelers. Applying booking limits to the fare products on a high demand flight can ensure that higher-fare passengers receive a seat instead of having the demand filled by low fare passengers. In the absence of booking limits, higher-priced fare product passengers would have a lower probability of securing a seat on the most popular flight departures. In addition to providing a superior level of service to the higher-fare passengers, this use of availability has economically efficient properties in terms of the resource allocation, which is formalized in the next chapter of this dissertation.

1.1.3 Pitfalls of Pricing Demand Management

A perfect passenger demand segmentation would be achieved at no cost to the airline and would have no effect on the demand for air travel in the lower-priced fare classes. The achievement of perfect fences as they are referred to in the industry is not an attainable goal for any airline because of the inherent imperfections in the passenger identification and segmentation process. The flexibility-based fences that prevail in the industry are imperfect in two main ways, through their inability to prevent passenger diversion to lower-priced fare products and in the decreased value of the lower-priced fare products resulting from the purchase restrictions.

The most often talked about departure from the perfect fence is the problem of passenger diversion from the higher-priced fare products to the lower-priced. Passenger diversion exists when passengers who would be willing to purchase higher-priced fare products purchase lower-priced fare products because they are able to meet the restrictions associated with the discount products. Passengers have incentive to divert from higher-priced fare products to lower-priced fare products when their value differential becomes less than the price differential between the two fare products. Fences that are unable to prevent passengers from diverting to lower-priced fare products are the fear of every airline and cause revenue dilution. For instance, American Airlines has estimated that only 6% of all air travelers purchased full-fare tickets in 1991 down from 35.6% ten years previous (WSJ, 1992).

The other problem encountered with the application of travel and purchase restrictions to the purchase of the discount fare products is the cost imposed by the segmentation mechanism. If the cost associated with accepting the increased bundle of purchase restrictions is greater than the utility associated with traveling on the flight departure at the targeted fare product level, an individual will opt not to travel. Airlines may circumvent the desired effect of demand stimulation if the cost perceived by consumers of the applied purchase restriction is too great. Considering the potential departures from perfect fences, the need for well designed fare products is apparent.

1.1.4 Fare Product Structure

Although there are countless airline fare products sold in many OD markets, a relatively small number of fare product types, or "selling fares", make up the greatest volume of sales. Among the primary selling fares there exists a hierarchy of airline fare products. In this hierarchy, the purchase restrictions facing consumers have travel and purchase requirements of increasing severity. The increasingly restricted fare products have progressively lower price levels, to avoid the case of a simple dominance of attributes. The more restricted fare products are also subject to increasingly reduced availability as a result of the seat inventory control techniques practiced.

The "value pricing" structure proposed by American Airlines in Spring of 1992 offered an example of an inferiority hierarchy of fare products. The only three

21

types of coach cabin fares available under the value pricing structure were the unrestricted fare product along with the seven and fourteen day excursion fares (E7NR and E14NR fare products, respectively). Fare products that require a minimum stay at the destination are called excursion fares. Today, most excursion fares are also non-refundable and have an advance purchase requirement. In this example, the unrestricted fare product clearly dominates either of the excursion fare products in all attributes except for price level since there are virtually no additional requirements for its purchase. In a similar fashion, the seven day advance purchase excursion fare product is preferable to the fourteen day. All product attributes of the seven and fourteen day excursion fares are identical except for the advance purchase requirement, fare level, and availability. The fourteen day excursion fare product is required to be lower in price level or face domination due to its inherent inferiority in other attributes.

Other non-excursion fares that have been offered in the past have been the three and seven day advance purchase fare products that require no Saturday night or other minimum stay. The seven day non-refundable excursion fare product has been sold previously with a fifty percent refundability penalty attached. The value pricing initiative sought to remove these fare products from the market to simplify the purchase and sale of fare products to consumers. The attempt failed, however, and American no longer attempted to defend the price levels in early October 1992. Even with the addition of the above fares, an unambiguous ordering of fare products by attributes can be achieved. The inferiority hierarchy of attributes and decreasing fare levels can be used to characterize the majority of fares sold in the market.

An example of the fare product structure existing in the market today is taken from the Portland, Maine to Dallas-Fort Worth, Texas (PWM-DFW) market. The selling fares in that market for travel on July 23, 1993, are listed in Table 1.1 as they appeared in the computer reservation system in early June. The fare product inferiority hierarchy is maintained in this example since the unrestricted fare product is the superior product on the market (independent of price level) followed in descending order of quality by the three, seven, and fourteen day excursion fares, respectively.

Fare Product	Price
Unrestricted	$1,120
E3NR	$600
E7NR	$530
E14NR	$470

Table 1.1 : PWM - DFW Summer Fare Product Structure - Early June, 1993

Table 1.2 shows the changes in the fare product structure with the introduction of a 30% off sale by Northwest Airlines in mid-June. The new seven day advance purchase fare made the earlier fourteen day and seven day excursion fares obsolete. The old $530 and $470 dollar fares are no longer selling fares in the PWM-DFW market because they are dominated by the new seven day excursion fare product. The new fourteen day advance purchase off-peak fare, valid for travel only on Tuesday, Wednesday, and Saturday, offers the new most highly restricted fare product on the market. It should be noted that the inferiority hierarchy is still maintained with the products ranking in descending order of non-price attribute quality from the unrestricted product to the three, seven, and fourteen day excursion fares, respectively.

Fare Product	Price
Unrestricted	$1,120
E3NR	$600
E7NR	$530
E14NR	$470
New E7NR	$353
Off-Peak E14NR	$329

Table 1.2 : PWM - DFW Summer Fare Product Structure - Mid-June, 1993

The availability of fare product purchases reinforces the fare product inferiority hierarchy. To demonstrate the increasing benefit of booking limits to the higher-priced fare products, the concept of a nested seat allocation is presented. A nested allocation scheme appears in Figure 1.1 for a three fare product example.

Nested Seat Allocation
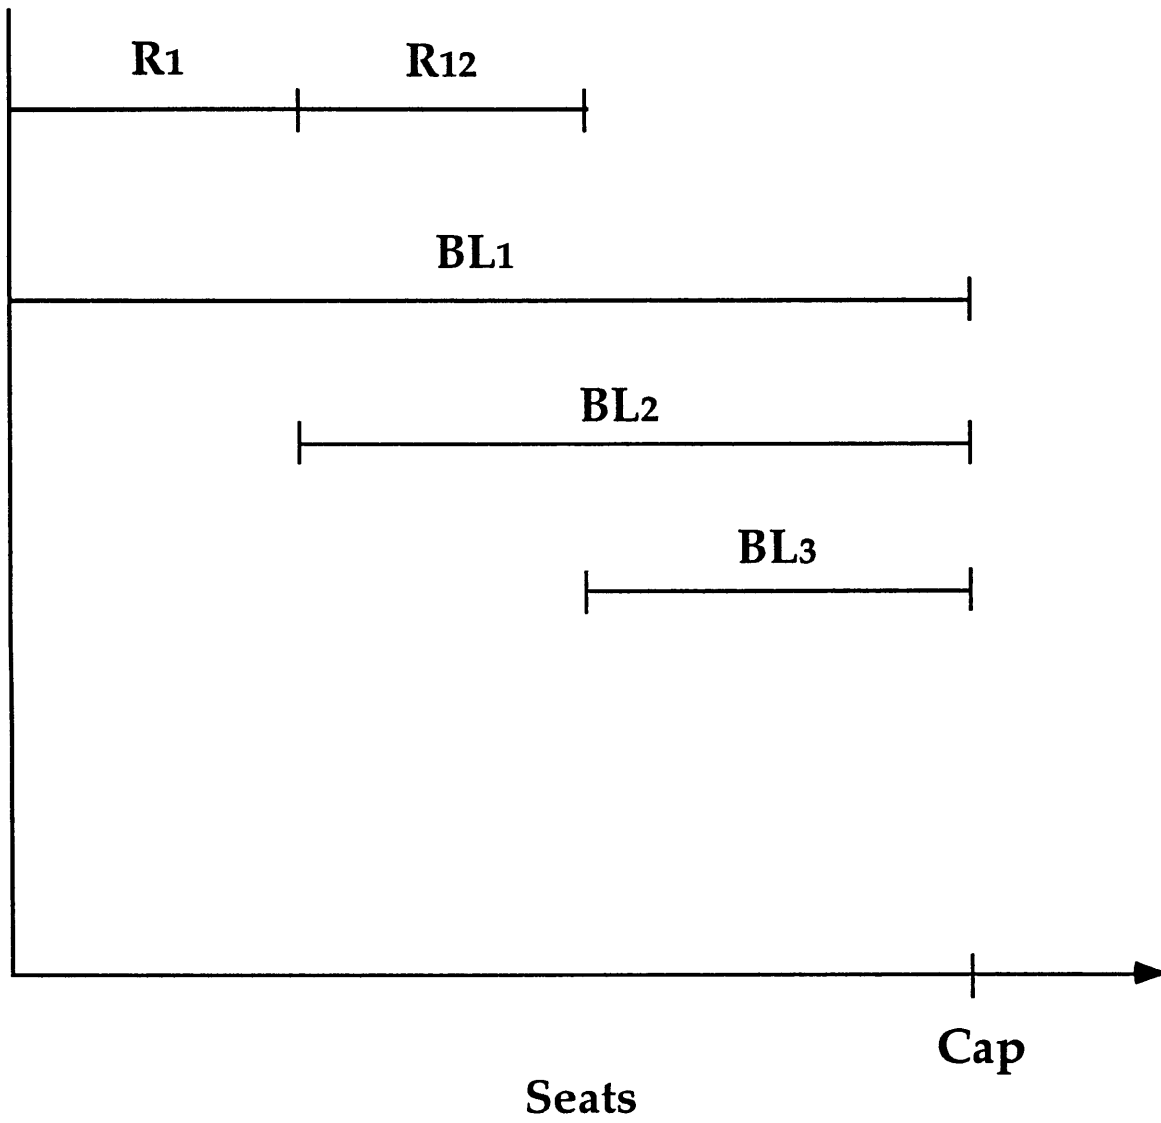
Figure 1.1 - Nested Seat Allocation

The fare products are decreasing in fare level and increasing in restrictions from fare product 1 to fare product 3. In the example, fare product 1 passengers have access to the entire aircraft capacity as their booking limits or $BL_1 = CAP$. Since fare product 1 passengers are paying the top price of anyone on the aircraft, the effect on revenue would be negative if a fare product 1 passenger requesting a seat were rejected (if one was available).

There are R_1 seats protected for fare product 1 passengers and the rest of the seats $(CAP - R_1)$ are available to fare product 2 passengers. R_{12} seats are protected for the use of either fare product 1 or fare product 2 passengers. Finally, fare product 3 passengers have $BL_3 = (CAP - R_1 - R_{12})$ seats which they are eligible to purchase. The seats reserved for fare product 3 passengers, however, can always be purchased by passengers from fare products 1 and 2 who value the service more (and are willing to pay for it). The nesting ensures that any unfilled seat is more likely to be available to the less-restricted fare products than the more-restricted thus reinforcing the inferiority hierarchy. The implementation of seat allocations that draw from a nested inventory dominate industry seat allocation structures.

1.2 Airline Revenue Management

The revenue management function within an airline controls much of the marketing of individual airline seats to consumers. The two major functions of the revenue management responsibility within an airline are pricing and seat inventory control (or yield management). The pricing department controls the airline fare product offerings and price levels in the market. The seat inventory control department allocates the seats given the fare levels and purchase restrictions set by the pricing department. All revenue management functions are undertaken assuming a fixed schedule and capacity provision in each market. The primary goal of revenue management is to manage the airline demand and revenue given the fixed provision decisions. A brief description of the revenue management functions appears in this section.

1.2.1 Airline Pricing

Pricing departments at airlines seek to provide airline fare products that stimulate demand where it is needed without unnecessarily diluting revenue from passengers placing greater value on air travel. In addition, the price levels and restriction bundles offered on the market must enable the seat inventory control department to best manage revenues given the fixed fare product structure. The published fare product structure for any airline must be flexible enough to manage demand under varying conditions, yet simple enough to be understood by travelers and ticketing agents. All of this must be achieved while remaining competitive with the fare product structures of rival carriers. In most cases, airlines offer identical fare product structures to their competitors as a result of the underlying market forces.

Other areas involved in the pricing function are the negotiated group pricing division and the negotiated contracts pricing division. The negotiated group pricing division bargains with tour contractors and representatives of large travel groups. The group negotiations are, in general, for a single flight or short term event. The negotiated contracts pricing function attempts to secure the exclusive business of a large client by offering a special deal for all employees or business trips on a long term contractual basis. The contract guarantees that all business from the single large client comes to the airline rather than going to other carriers. The contract, in general, fixes a corporate rate for a large purchaser of air travel that is guaranteed for the duration of the agreement. It should be noted that the negotiated group and contract pricing decisions must both be made while competing for seat availability with passengers traveling on published fares.

The focus of this dissertation is the establishment of price levels and fare products that enable an airline to maximize revenues. It is therefore logical that the focus of the discussion of an airline pricing department be the published fares and fare products. A detailed description of the daily published pricing function appears later in the chapter.

26

1.2.2 Airline Seat Inventory Control

The airline seat inventory control function within an airline seeks to maximize revenue given the fixed capacity flying between city pairs, the fixed price levels and restrictions facing them (set by the pricing department) by forecasting passenger demand and correctly allocating the available seats. To this end, airline seat inventory management seeks to reduce revenue dilution and prevent unnecessary decreases in revenue and yield. Analysts control the availability of seats using the inventory controls available to them in the computer reservation system (CRS).

Computer reservations systems allow airlines to control their seat inventories virtually in real-time by recording and communicating passenger bookings on the large computer network. The inventory control ability has enabled the airlines to increase their revenues given the prevailing fare levels and demand forecasts using the techniques of yield management. The three main areas of yield management are demand forecasting to predict traffic volumes, booking limit applications that attempt to maximize expected revenues, and flight overbooking to correct for the problem of passenger "no shows". The sophistication of airline yield management modeling has been implementable because of the CRS network.

The seat inventory management department applies booking limits to the available fare products to maximize the expected revenue of the airline. The application of booking limits is the function most closely associated with the seat inventory management department. To maximize expected revenues on the flight departures, booking limits are set using the techniques of yield management such as the Expected Marginal Seat Revenue (EMSR) heuristic (Belobaba, 1987). Yield management systems use demand forecasts as inputs to maximize expected revenues.

Demand forecasts must be made by the seat inventory management department to predict the number of passengers that are expected to make fare product bookings on a given flight departure. The passenger traffic predictions are a primary input to the yield management systems which enable the airline to best allocate the available seat inventory to maximize expected revenues. Without knowledge of the number of passengers expected to book on a given flight

departure, the airline would have no idea how many seats to protect for the high revenue passengers. Thus, the ability to apply sophisticated demand management techniques hinges upon the ability to correctly forecast demand.

Finally, not all passengers who make bookings on a flight departure actually show up and depart on the flight. The result of the passengers who do not show up is the departure of the aircraft with empty seats. If there is excess demand on a given flight, allowing a passenger to book a flight and then cancel results in lost revenue for the airline as well as a wasted resource. The seat inventory control department overbooks flights, accepts more bookings than the aircraft capacity, to prevent seats from departing unfilled. Setting overbooking factors for flights is another primary function of the seat inventory management department. The overbooking factors set for a flight also impact the seat allocation decisions returned by the yield management system.

1.2.3 Cooperation Between Seat Inventory Control and Pricing

Although there is a dependency relationship between pricing and seat inventory control, the policies set by pricing must be made first. The pricing department sets price levels and the accompanying restrictions to the fares offered in each OD market. Seat inventory control takes the fare levels and restrictions established by pricing and attempts to maximize revenue. All actions taken by seat inventory control are contingent upon the price levels and fare restrictions dictated by pricing. It is the responsibility of pricing to provide the price levels and fare products required by seat inventory control to correctly manage demand and revenue.

Within the pricing department, the correct price levels and fare restrictions must be combined to offer a competitive mix of fare products. Seat inventory control is responsible for managing revenue and does so by protecting seats for higher-fare passengers booking near departure and limiting availability of discount seats. When pricing must issue a competitive price filing to match a newly introduced lower fare, seat inventory control is responsible for applying booking limits to the number of fare products that may be reserved by consumers at the published price. A miscommunication between pricing and seat inventory control could result in too many discount fares to be sold on a given flight leaving the profitability of the flight diluted. Proper communication between

28

pricing and seat inventory control can help to prevent unnecessary revenue loss for the airline.

Drastic changes in price levels may impact the seat inventory control function in predictable ways. The demand forecasts may become biased with large shifts in price levels. For example, during a fare war, the number of bookings in the lowest discount class are expected to balloon as a result of the low fares being published. The seat inventory control analyst would be better equipped to anticipate demand spikes with proper communication with pricing. Any decisions made by pricing and inventory management should consider the possible effects that the pending decision might have on the revenue dilution encountered when high-revenue passengers are displaced by low-revenue passengers. Although generally regarded as separate within airlines, the pricing and seat inventory control functions work to achieve the same goals of demand and revenue management. The ultimate aim of the revenue management function within an airline would be a complete integration of the pricing and seat inventory control functions.

1.3 Current Practice : Airline Pricing Departments

The primary functions of airline pricing departments are to develop the differentiated fare products, set their price levels, and report this information to the relevant systems. This section seeks to characterize the daily functions and requirements of the published fares portion of an airline pricing department. A description of the communication of price levels between airlines begins the discussion.

1.3.1 Airline Tariff Publishing Corporation

In the U.S. airline industry, each airline knows the complete fare structure of all other carriers in every OD market. Domestically, the fare product information is contained in a data base published by the Airline Tariff Publishing Corporation (ATPCO). Airline fare product changes are sent at least daily from the individual carriers to the ATPCO. The Airline Tariff Publishing Corporation serves as a clearinghouse for the fare product changes that are posted by the carriers.

The information concerning fare products available in any OD market that is held by the ATPCO data base includes :

1) Fare level
2) Fare basis (a code representing restrictions)
3) Effective dates
4) Acceptable routings
5) Special footnotes
6) Specific rules applicable
7) Airline quoting the fare

The fare level posted with the ATPCO is simply the cost to consumers of purchasing the fare product. The fare basis is a code representing the restrictions associated with a fare product and the fare class in which it is published. For example, a QE14NR fare basis code identifies a fare product booked in Q class for an Excursion fare that must be purchased at least 14 days in advance of departure and is Non-Refundable. The effective dates comprise the first and last dates that the fare is available for purchase by the traveling public. Any specific restrictions on the flights for which a given fare product is valid are contained in the acceptable routings category. Special footnotes and rules clarify any non-standard restrictions placed on the fare product or its sale not evident in the fare basis code or the routing restrictions. Changes can be posted with the ATPCO by the carriers to any of the above information fields, not just price level. From the ATPCO data base, the current airline fare product structure for all carriers in all markets is known for the present and near future to all subscribing carriers.

The fare levels and products listed with ATPCO that are currently valid (effective) are also listed in the computer reservations systems that provide information about the inventory of seats for sale to the public. The fare level changes are also reported to the CRS network on a frequent (at least daily) basis. The fares listed are the only ones that can be sold through the CRS. As previously stated, the unpublished negotiated and contract fares must be handled separately by the airline and do not appear within the ATPCO data base or any CRS. The primary differences between the CRS and ATPCO data bases are that the CRS contains seat availability information while the ATPCO contains a greater volume of information about fares not currently available for sale to consumers.

The fare changes that have been registered with the ATPCO data base in the last 24 hours are published daily in a report made available to all subscribing airlines. Using the data base, responses to price and fare product changes can be registered and disseminated within 24 hours. In this framework, discontent with a fare change can be made known to another airline within one day. As would be expected, the ATPCO data base has been used by the airlines as a method to engage in potentially collusive pricing practices, most notably price signaling.

1.3.2 Price Signaling

The practice of price signaling using the ATPCO published fares data base is one technique available to the airline pricing analyst. Price signaling is the method by which airlines show discontent with the price changes of other carriers. The airlines test the competitive reaction to proposed fare initiatives before they go on the market using the ATPCO clearinghouse. Carriers "signal" their intention to introduce a fare product change or price reduction onto the market in advance, to see the competitive response of the rival carriers before actually putting the fare product on sale. When price signaling is used, consumers need not receive the benefits of low fare quotes that are not accepted by the other carriers since the changes need never reach the market. The acceptability to the competing carriers is often the force that drives the actual introduction onto (or withdrawal from) the market of a fare initiative.

For exposition, consider the example of a fare decrease in the Chicago to Dallas-Fort Worth market posted by United Airlines to take effect in one week. If American Airlines disapproved of the pricing action by United, American might propose a fare increase in the same market identifying United in the fare basis code with effective dates for the fare identical to United's proposed fare reduction. From this signal, United would understand that American disapproved of the price reduction in the market. A more severe reaction to United's fare cut by American might be to publish an extremely low fare in the Chicago to Denver market (which is very important to United) effective for only one day. This would indicated extreme displeasure to the fare change by American. It would be up to the discretion of United whether or not to withdraw the initial fare decrease, but that decision would be aided by the information about the competitive reaction of American as shown through its price signals.

The message sent by a carrier can vary in severity from a minor fare modification in a low revenue market to a large fare cut in an important hub market. If a carrier views the price level in a certain market to be below a level at which a profit can be made, the carrier may send pricing signals in an attempt to raise the prices up to an acceptable level. To be most effective, this process must be done gradually and in systematically timed increments so that the signals can be clearly interpreted by the competitors (Elkins, 1986). To effectively price signal, the other carriers must know your intentions. Regardless of the method of implementation, however, the practice of price signaling pervades the industry.

Legal and ethical issues surround the current practice of pricing in the airline industry. Among the volatile issues facing airline management is the legality of the practice of price signaling. One of the major concerns facing the United States Department of Justice regarding the commercial airline industry is the legality of price signaling and when signaling becomes a violation of anti-trust legislation. The U.S. government has disagreed with the testing of fares before they go to market. A lawsuit has been filed by the U.S. government against eight major U.S. carriers in an attempt to curb price signaling. The carriers were found guilty of anti-trust violations and ordered to desist from the practice of price signaling. As a result of a settlement, two major carriers, United and USAir, have signed an affidavit preventing them from practicing fare testing using the ATPCO clearinghouse for the period of ten years (WSJ, 1993).

1.3.3 Functions of an Airline Pricing Department

Airline pricing departments are made up of analysts who make the daily pricing decisions and managers who approve the decisions and handle more far-reaching pricing issues. Analysts are assigned several OD markets for which they are responsible. Primary in the duties of the analyst is addressing the fare actions of other carriers in their markets. Fare initiatives also must be registered by the analysts at the request of pricing management. Figure 1.2 represents the flow of information within the pricing department and out to the ATPCO that occurs daily in all OD markets where fare change actions take place. Changes to restrictions and other fare product attributes also take place in the ATPCO clearinghouse. These changes are quite similar in nature, however, and thus, the discussion will focus on fare changes.

Overview of an Automated Airline Pricing System
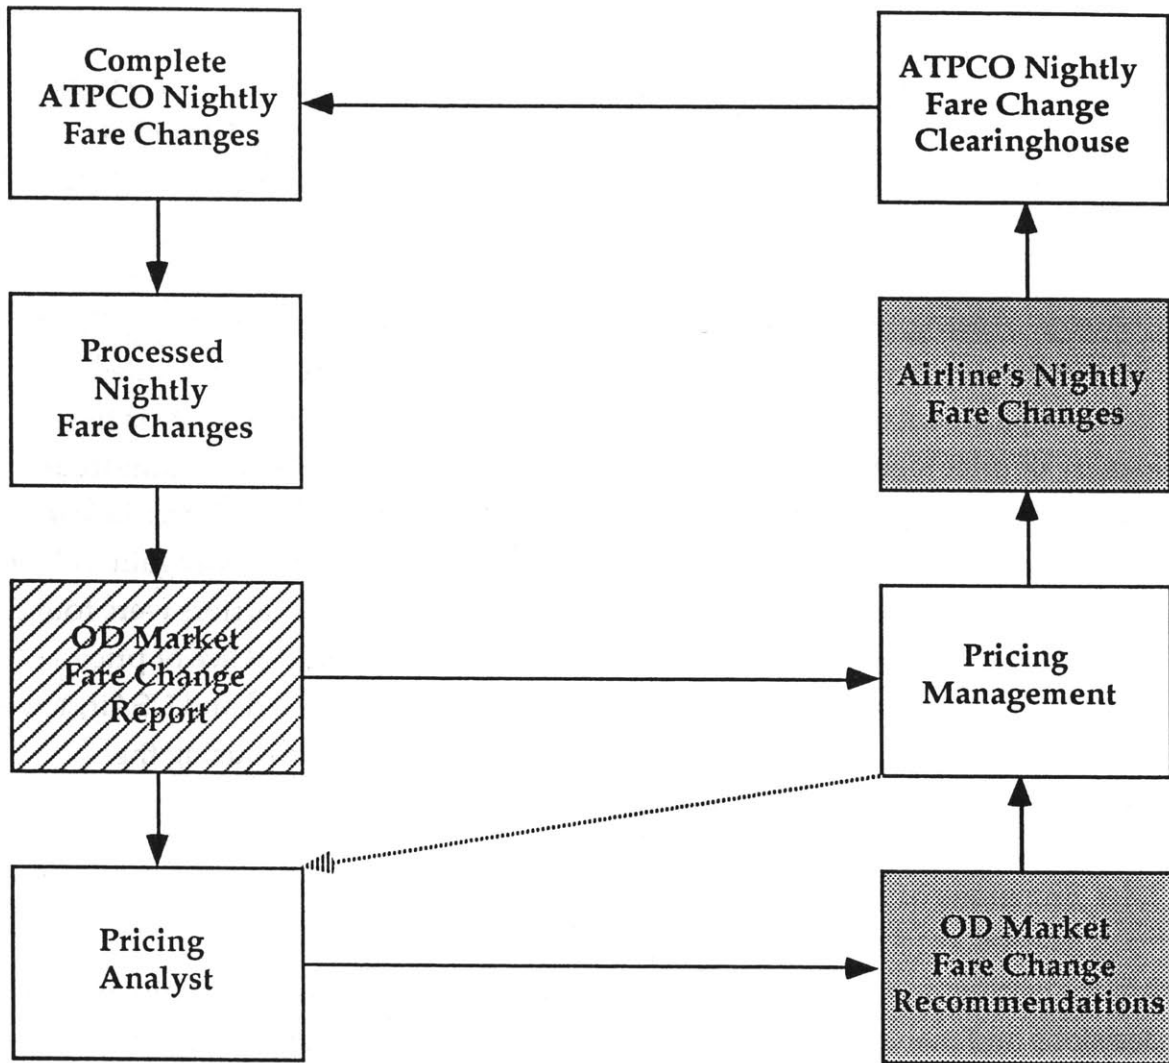
Figure 1.2 - Overview of an Automated Airline Pricing System

Each day the ATPCO returns a list of the fare changes that were posted the previous day by all of the participating carriers. In Figure 1.2, the cross-hatched box represents the information provided to the analyst about his OD markets. The lists of changes are processed by the carriers and provided to pricing analysts in the form of market reports. In the past, the reports had been distributed to analysts in the form of computer printouts on their markets. Many major carriers now process this information and provide it to analysts on menu-driven screens. An efficient information processing function can expedite a thorough market analysis for the pricing analyst.

Analysts access fare changes in the markets for which they are responsible to monitor the activity that has taken place over the past day. The shaded boxes in Figure 1.2 represent the changes made by the analysts for submission to the ATPCO clearinghouse. It is the responsibility of each analyst to respond to the fare changes posted by rival carriers in his markets. The fare change responses or initiatives that are proposed by analysts are then transferred to pricing management for approval or revisions. Upon approval, the pricing decisions that have been made in all markets are sent to the ATPCO data base for listing in the fare change report provided to all carriers the following day. It should be noted that communication between the airline pricing departments and ATPCO may take place more often than once per day but for ease of exposition, a daily interaction has been assumed.

1.3.4 Competitive Pricing Strategies

When a fare change is registered by an airline, a competitive response is expected from the other airlines offering service in the OD market. This is true whether or not price signaling is permissible. The discussion of competitive responses to fare changes is presented here with respect to a decrease in fare level although the result can be generalized to any change in fare levels or purchase restrictions. Three main types of price quoting strategies are used to respond to a fare change when registered by a competing carrier:

1) matching strategy
2) non-matching strategy
3) partial matching strategy

If the responding carrier agrees with the price change or is unwilling or unable to combat the fare change, a matching strategy can be adopted. A matching strategy consists of matching the price level offered by the carrier posting the fare change under the same conditions on all flights in that market. The expected results of a matching strategy are the retention of market share and competitiveness in the market. A negative effect resulting from a matching strategy (in this case) is a loss in the average revenue paid per passenger in the market because of the reduced price levels.

In contrast to the matching strategy, a carrier can adopt a non-matching strategy in which the carrier retains his current price levels on all flights. This strategy would result in the maintenance of yield but risk the adverse effects of a possible loss in market share and competitiveness in the market. In addition, customer goodwill may be lowered in the market if the carrier is perceived as being insensitive to the traveling population by not matching, for example, a heavily advertised fare cut. When both the matching and non-matching strategies offer unacceptable alternatives, any of several partial matching strategies which call for selective matching of fare changes can be employed.

A spectrum of partial matching strategies lie between the matching and non-matching strategies which may help in retaining market share and reducing the loss of yield. A partial matching strategy may result in a price match on certain off-peak flights or on flights with a specific routing. For example, one partial matching strategy calls for the matching of fares on the less popular, off-peak flights while maintaining current fare levels on the more popular peak flights. Another partial matching strategy calls for routing restrictions on the discount fare products in an attempt to reduce the pressure put on high load factor flight legs entering and leaving hubs burdened with too much rejected demand.

The exact form of the strategy adopted by the carrier should depend on the nature of the OD market and the relative strength of the carrier in that market as well as the size of the fare change. The ramifications of adopting a partial matching strategy must be considered when choosing the best fare-matching response strategy. The competitor offering the lower fare is likely to offer that fare on all flights (peak and off-peak) in that OD market. The pricing analyst must be confident enough in the market demand level at peak hours to become

non-competitive on price level for the peak flights. The experience and judgment of the analyst is at a premium when evaluating partial matching strategies.

The carrier instituting fare changes must carefully monitor the pricing reactions of the other carriers to gauge their relative agreement or disagreement with the price level change. While the primary method of fare change response is registered in the affected market, competing carriers disapproving of a fare initiative may display their discontent in other markets. Competitive responses to fare changes are often a system-wide phenomenon. Thus, communication of price changes between pricing analysts is another important function of the pricing department.

1.3.5 Monitoring Fare Changes Properly : Practical Considerations

The direct and indirect effects of having discrepancies between prices of like fare products must be examined when fare differentials result from strategies other than direct matching. When price differentials are enacted in a given market, the effects should be monitored closely to gauge the impact on revenue, traffic, and yield. Yield is the revenue per passenger-mile earned by a carrier. If the number of bookings in the affected market is substantially lower than before the activation of the fare differential, it is likely that the differential should be removed from the fare listings. It is important to monitor both load factor and yield when price level changes are being instituted to gauge the resulting effect on revenue.

When a fare differential is listed in a market among the ATPCO published fares, it is the responsibility of the pricing analyst to verify that the fare level is actually available to consumers. Not all fares appearing in the ATPCO data base are actually available to consumers as a result of, for example, no service provision by a carrier posting a fare. Or, there may be no authorized seat availability at a certain fare level appearing in the ATPCO data base. It would not be logical, for instance, for a dominant carrier to match a lower fare quote by a carrier offering virtually no seats in a market if the dominant carrier did not approve of the lower fare.

Another particularly detrimental effect of fare differentials is the removal of the quoted fare from the first page of the CRS screen display of like fare products.

For instance, when a potential customer calls a travel agent and asks for the lowest fare available in the market, the agent is likely to view the lowest priced fare listing on the first page of the price-stratified menu display. Entries appearing on the following pages are less likely to attract the attention of the travel agent and are thus, booked less frequently. Publishing a fare differential on the lowest fare in a market may result in the removal of that fare from the first page of the CRS display. Even if the higher fare is not removed from the initial page, it will appear after the lower priced ones on the display. The convenient availability of the fares to travel agents via the CRS display is a critical consideration for the pricing department that is not outwardly apparent.

Finally, carriers choosing to follow non-matching or partial matching strategies must be aware that many passengers choose an airline and a specific flight on the basis of price alone. Such increases, if unmatched, would potentially result in a decrease in load factor and revenue because of the many passengers who are so highly price sensitive. As a result, unilateral fare increases in a single market are not a common practice at most airlines.

1.4 The Nature of Airline Pricing Decisions

The price levels facing consumers in major U.S. markets rarely differ by carrier. The fear of "fare wars" prevents the major carriers from competing on price level. Moreover, the rapid communication of pricing information results in identical fare products being posted in a market at almost all times. The method by which the "equilibrium" of fare products is reached, however, merits discussion. Many different factors influence the pricing decisions that are made in a market. Among the most important are the level of service offered by the competing carriers and the relative strength of the carriers in the market. Perhaps the most important level of service considerations are the existence and frequency of non-stop service in the market. To the price-insensitive traveler, the existence of non-stop service is an extremely important level of service consideration. This section seeks to address price level equilibria and the factors that influence them.

1.4.1 Fare Product Equilibria : Price Levels and Restrictions

The prevailing prices in any OD market have most likely been arrived at through some sort of competitive equilibrium pricing mechanism. The equilibrium achieved is not, in general, the one or two price level type assumed to exist in many models of airline competition. The price levels and fare product types are generally matched across all carriers and all fare products based on market conditions. Price levels tend to reach a level of competitive equilibrium where the fare and restriction combinations offered by every carrier are almost identical. The competitive price level equilibrium forces competition to take place primarily on the non-price level. The equilibrium prices of fare products arrived at in the market will most likely depend upon the number and relative strengths of the competitors in the market. Other important factors include the level of service offered, the demand for service in the market and, to a certain extent, the cost of providing service.

Carriers offering non-stop service or greater departure frequency can exercise more influence over the prevailing market price levels than other carriers as a result of their greater market power. In terms of fare product types, the basic industry structure (or a close variant) is likely to rule the market structure. Variations from this may be driven by the peaking behavior of the market. For example, large day of week or time of day differentials in demand may lead to the introduction of a peak/off-peak pricing structure in certain markets.

The rule of thumb concerning fare products existing in the market is usually that carriers tend to match fare structures by market. The incidence of non-matching and partial matching strategies is quite limited. A competitive equilibrium of price levels and fare product restrictions is the rule rather than the exception in day-to-day airline pricing decisions. A single fare level associated with each fare product is thus a reasonable assumption when modeling airline pricing practice with its resulting fare product structure. In general, the competing airlines appear to act as one entity when the final price levels prevail on the market. Thus, although the prevailing price levels have resulted from an equilibrium process, the result appears to consumers as a single price level for each of the different fare products.

Unlike most industries, when airline pricing analysts price fare products in a given market, operating costs are not considered explicitly. Instead, pricing analysts look at the prices that currently prevail in the market. Despite their absence from daily pricing decisions, however, the operating costs for the flights may be referenced to determine whether a profit can be made at the established price levels. If after analyzing the profitability of a given market there is found to be an operating loss, a change in the type of service offered is considered or an attempt is made using price signaling to raise the fares. In the long run, service is likely to be discontinued in that market if the operating loss persists. Such decisions are not within the standard operating jurisdiction of the pricing department and are addressed by the scheduling department or top level management.

1.4.2 Hub and Spoke Network Effects on Pricing Decisions

Since the enactment of the Federal Airline Deregulation Act of 1978, there has been a trend by U. S. airlines to tailor their route structures into a series of hub and spoke networks. The hub and spoke structure takes advantage of the "network effects" gained by routing many flights into a single hub and allowing passengers to connect to any of the other destinations served by the hub. A connection between almost any OD market pair can be accomplished by making no more than two connecting flights when a hub and spoke network structure as vast as that in the U.S. domestic air transportation system exists.

A hub and spoke structure allows the airline to serve far more OD pairs than would be possible by providing point-to-point service between the OD pairs. In the context of total markets served, the hub and spoke network configuration used by all major airlines today is clearly a more efficient use of resources than would be the case if the same amount of resources were deployed to run a point-to-point network system. Consider the example of a five city point-to-point route network shown in Figure 1.3 which requires twenty flights per day to offer non-stop service between the five cities. Next, consider the five city hub and spoke network shown in Figure 1.4 requiring only eight flights per day to offer direct one-stop service between all OD pairs.

Point-to-Point Airline Route Network
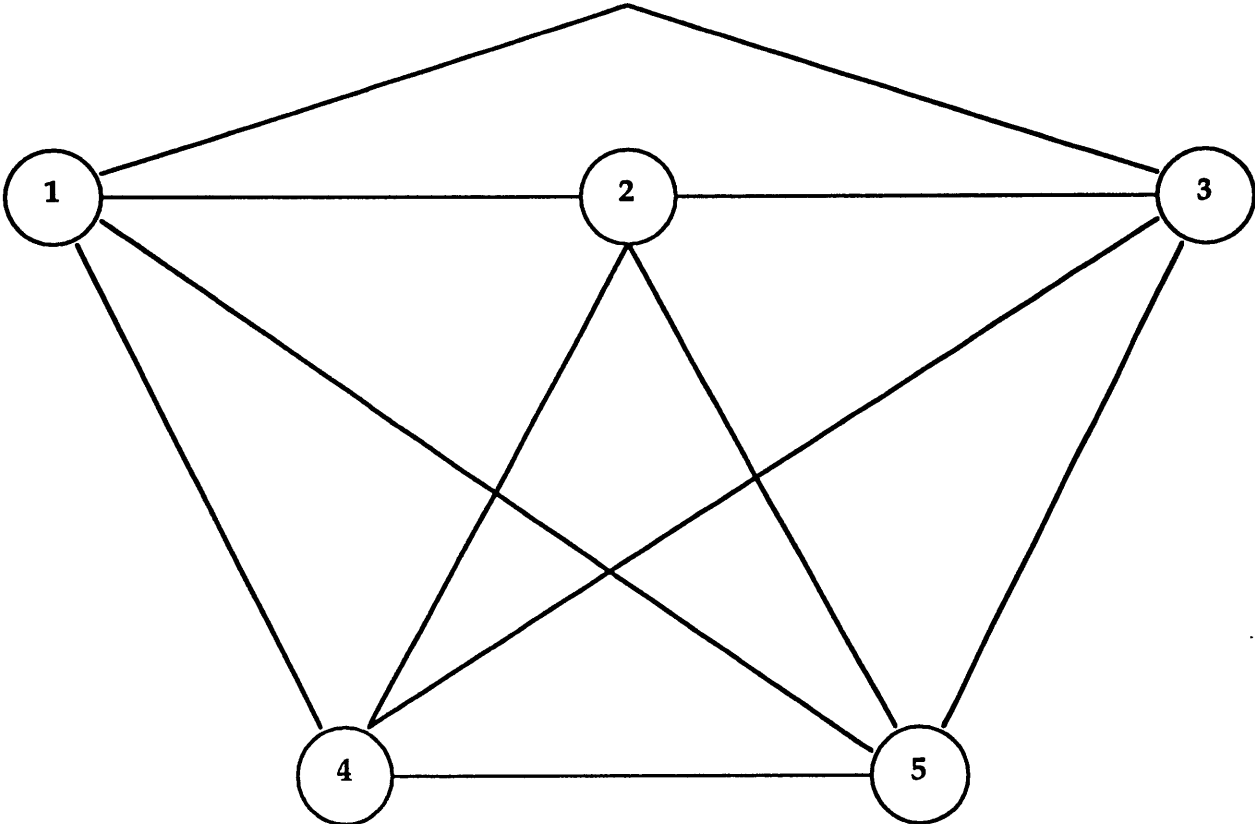
Point-to-point airline route network serving 5 cities with 20 links which are presented as 10 bidirectional flight legs

Figure 1.3 - Point-to-point Airline Route Network

Hub and Spoke Airline Route Network
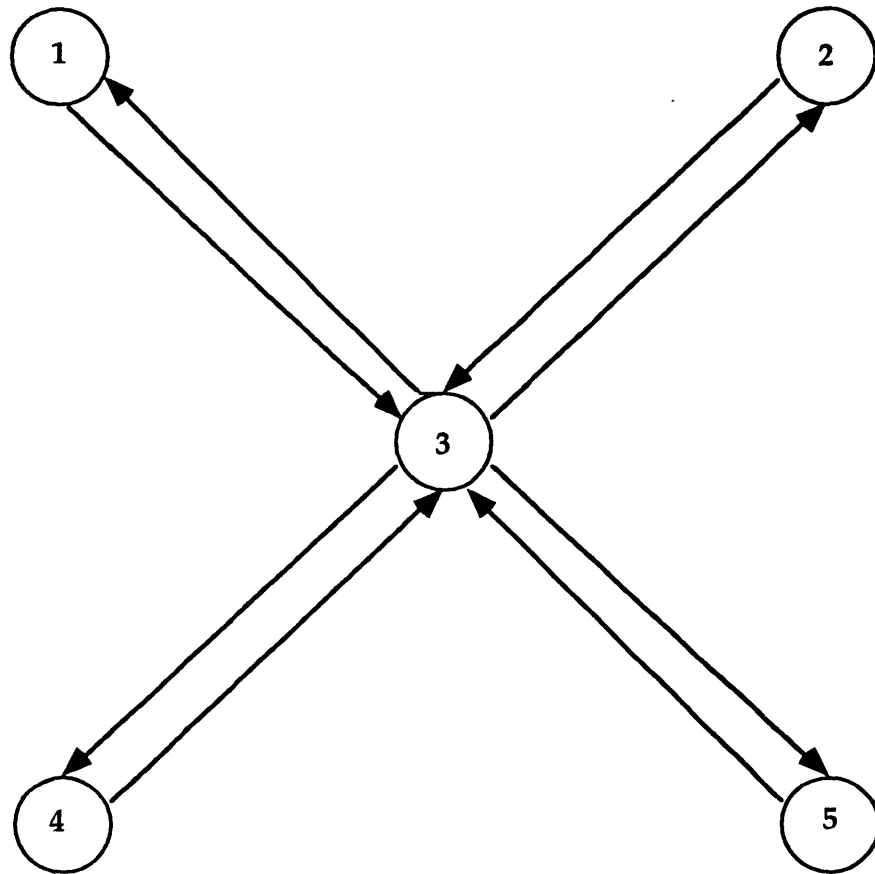
Hub and spoke airline route network configuration
serving 5 cities using only 8 flight legs

Figure 1.4 - Hub and Spoke Airline Route Network

The simple example demonstrates the benefits in resource allocation offered by a hub-spoke system which increase with network size. The number of flight legs required to serve a point-to-point network amplifies factorially with increases in the number of cities served by the network while in a hub-spoke network it increases by only two. Thus, the larger the size of the network (*e.g.* the larger the scope of operations of an airline), the greater the number of services that can be offered to a larger number of people. Each new spoke added to an existing hub-spoke network can generate an array of destinations available to all locations served in the network requiring only a single stop at the hub. Therefore, despite the slight decrease in the level of service offered by the hub-spoke configuration, the ability of the hub-spoke system to capitalize on network effects makes it cost-effective compared to a point-to-point network.

The result of the economies of scope prevailing in the industry has been a move by most of the major airlines to offer their services as competing hub and spoke networks. Airlines not using the hub and spoke may be forced out of the market by the relative inefficiencies encountered when offering a point-to-point service. This view is based solely on cost, however, and the influence of demand levels and passenger preference cannot be ignored. Since virtually all passengers prefer non-stop service to direct or connecting service, the carriers offering the non-stops have an advantage on the demand side, both in passenger attraction as well as market pricing.

Any airline operating a point-to-point service system would be unable to compete with the established hub-spoke systems all things being equal unless that airline can operate under an extremely low cost structure. Factors other than cost, however, often influence the ability of a carrier to compete more greatly. The example of Southwest Airlines deserves mention since they operate a primarily point-to-point route network quite profitably in the United States. The ability of Southwest to compete stems from low operating costs, a monopoly at Love Field in Dallas, high employee productivity, and passenger preference for non-stop service. The markets served by Southwest are, in general, short distances in which non-stop service is all but required. Thus, although cost-effective, the efficiencies offered by a hub and spoke are by no means a necessary ingredient for success.

42

In any case, the type of competition occurring in the airline industry is influenced by the hub and spoke route structure that prevails at almost every major North American carrier. Because of the hub and spoke network configuration, many carriers are able to provide service in the largest OD markets. In turn, the increased number of OD markets that can be served by carriers increases the level of competition in most medium to long haul markets.

The competition invites attempts at price competition by the weaker carriers. For instance, carriers in need of cash to sustain their operations often attempt to increase market share and loads by offering price reductions to passengers. Since stronger carriers are adverse to allowing fare differentials in most markets, the result is lower fare levels across the board due to matching. The discounts may even be below levels at which a profitable return on revenue can be achieved. The desire of carriers to offer a full range of air travel services along with the costs of entering and exiting markets frequently forces carriers to accept unprofitable price levels in the short term. The wide range of OD markets that are served by weaker carriers produces a downward influence on price level since, in most cases, all carriers match the lowest available fares.

The existence of the hub and spoke air transportation network also allows carriers to offer increased frequency of connecting service compared to what can be offered in a point-to-point route network with an identical aircraft fleet. An important consideration to the price-insensitive traveler is the number of flights offered per day by each of the airlines in the OD market. Because the fare products purchased by price-insensitive travelers generally allow itinerary changes, it benefits these travelers to choose the airline that offers the most frequency of service for the OD market in which they are traveling. Carriers offering the most frequency in a market (and particularly the most non-stop frequency) are often the most powerful carriers in the market. The dominant carriers are able to exercise a degree of market power when faced with fare changes. In this way, greater frequency of service affords carriers influence in market pricing decisions.

The variability in level of service across OD markets between the carriers with competing hubs changes the nature of competition. The level of service offered in markets varies widely by market and by carrier within the market. One result of the hub and spoke system is the difference in level of service offered in the hub

markets as compared to the non-hub markets. For instance, consider the Boston to Atlanta OD market. Only Delta and Trans World Airlines offer non-stop service. While other carriers offer service in the market, USAir through Charlotte for example, the pricing decisions are dominated by the non-stop carriers, particularly Delta which offers significantly more frequent non-stop service than TWA. USAir, on the other hand, dominates the Boston to Charlotte market as a result of its large hub in Charlotte and strong presence in Boston. Clearly, the competitive stance of every airline differs from market to market as a result of route structure. From the perspective of the individual airline, the nature of competition differs in every market served and must be evaluated individually.

The position of each competing airline is different in every market. Market power weighs heavily in the determination of the prevailing price level of fare products in an OD market. Price levels in an OD market are generally most heavily influenced by the carrier offering the most non-stop frequency in the market. The most effective competitive measures are the existence of non-stop service and frequency of service. The perception of reliability of service remains a necessary but not a sufficient condition for competing in the market on the non-price level. If a carrier is not perceived as reliable by price-insensitive consumers, all other non-price level competitive moves are irrelevant. The level of service and perceived level of service offered by every carrier operating in the market must be examined when characterizing the factors influencing pricing power.

1.5 Outline of the Dissertation

The dissertation begins by formalizing the similarities and differences between airline fare product differentiation and price discrimination in airline markets in Chapter Two. The question of the economic efficiency of airline pricing is addressed in the context of the current environment of fare product differentiation and seat inventory control. The existing work done on airline price discrimination and fare product differentiation in both the economics and airline yield management literature is reviewed.

Chapter Three looks at the passenger arrival and booking process in detail. A model is proposed that addresses the shortcomings of the existing characterization of passenger demand in all yield management seat allocation

44

models. The independence of the demand for the individual fare products is relaxed in the model to more correctly characterize the situation facing airlines. The data requirements for such a model are, however, prohibitive and thus a more tractable model is sought.

Chapter Four introduces the generalized cost model for airline fare product differentiation, an operational static model that provides a more correct treatment of airline fare product demand. The model incorporates the dependence of fare product demand levels to more realistically characterize passenger demand. The model explicitly considers the decrease in product quality associated with accepting a more-restricted fare product by imposing a cost on each consumer purchasing a discount fare product. Functional specifications of the model are then presented.

Chapter Four continues by reviewing the modeling of passenger "buy down" or diversion to lower-priced fare products as it has been addressed in the literature. Passenger diversion is then incorporated into the generalized cost modeling framework. Finally, booking limits are applied to the static model to approximate the effects of yield management on the airline fare product differentiation model.

Chapter Five demonstrates the uses of the generalized cost model under the assumption of two types of functional forms of passenger demand. Both linear and constant elasticity demand specifications are incorporated into the modeling framework. The optimality conditions and limiting cases are highlighted to demonstrate the effects of the model. Practical applications of the generalized cost model are then presented. Finally, Chapter Six presents the conclusions and contributions of the research and provides a discussion of the future research directions.

1.6 Chapter Summary

The chapter began by reviewing the historical view of airline pricing and its evolution to the current structure. Airline pricing was discussed in the context of the fully deregulated environment. The revenue management function of airlines was then presented. Current practice within an airline pricing

department was detailed. The nature of airline pricing decisions was then discussed. Finally, an outline of the dissertation was presented. The context of airline fare product differentiation has thus been set so that a discussion of the methodologies used to interpret pricing decisions can now be undertaken. The next chapter discusses the view of airline fare product differentiation in the literature and in society.

Chapter Two

Airline Fare Product Differentiation Literature Review

Airlines might choose to describe the offering of selected air transportation and restriction bundles as product differentiation rather than price discrimination. Airline travel consumer advocates, on the other hand, might choose the opposite characterization. The two terms have been used almost interchangeably in the airline yield management literature leading to confusion about these two closely related, but distinctly different, concepts. Distinguishing between product differentiation and price discrimination is necessary in order to begin a proper review of the airline fare product differentiation literature.

Product differentiation is the practice of offering products with differing attributes for sale to consumers. The varied airline fare product offerings available on the market today are a clear example of product differentiation. Although the level of in-flight service offered to passengers in the coach cabin is identical, the travel flexibility associated with the individual tickets often varies widely. Each fare product offered by the airlines represents a unique collection of travel attributes.

In contrast, price discrimination is the practice of charging consumers differential mark-ups over the marginal cost of serving them. Price discrimination occurs if the price differential between higher- and lower-priced fare products differs from the marginal cost differential of serving the higher- and lower-priced fare product passengers. Thus, if it is no more expensive to serve the higher priced fare product passengers than those purchasing the lower priced, price discrimination exists. Such differentials indicate the existence of price discrimination, according to the strict definition found in the economics literature.

Under these definitions, product differentiation and price discrimination can occur simultaneously or separately. An airline need not offer differentiated fare products in order to price discriminate. Conversely, by offering differentiated fare products, an airline is not obligated to price discriminate. Since it is evident

that airlines practice product differentiation in virtually every OD market that they serve, it remains to establish the extent to which price discrimination is practiced.

This chapter begins with a discussion of the nature of airline price discrimination and fare product differentiation in airline markets. Airline fare product differentiation is then explained within the context of the economics price discrimination literature. Efficiency considerations are then addressed with respect to airline realities. Next, the existing models of airline price discrimination and fare product differentiation that have appeared in both the economics and airline yield management literature are reviewed. The assumptions and findings of the existing models are documented and the needs for future research are motivated.

2.1 The Nature of Price Discrimination in Airline Markets

Price discrimination occurs when consumers are charged different price levels with no differences in the marginal production costs of serving them. Most costs incurred by the airline do not differ between consumers, or consumer types, since it costs the airline nothing to add restrictions to the fare products. One notable exception is the opportunity cost incurred by allowing passenger to unconditionally return unused unrestricted tickets. If marginal costs with respect to passengers are constant, as is often assumed, the existence of fare differentials between different consumer groups results in discriminatory pricing. In this case, airlines unquestionably practice price discrimination since the price level mark-up over marginal cost is greater for one population than another.

In the economics literature, price discrimination has been characterized as occurring in three forms, first, second, and third degree, based upon the seminal work of Pigou (1920). A brief description of each type of price discrimination follows with its specific application to airline fare product differentiation discussed thereafter. Only a cursory treatment of price discrimination techniques is presented here in order to classify airline fare product differentiation and price discrimination within the framework of the economics literature. For a more

48

detailed discussion of the different types of price discrimination consult Tirole, 1988.

2.1.1 First Degree Price Discrimination

The practically unrealistic case of first degree price discrimination assumes that the airline is able to perfectly identify and segment each potential consumer and force him to pay a value equal to his maximum willingness to pay for the product. Such a segmentation would result in the maximum profit for the airline with no consumer surplus. Even with the ability to achieve such a perfect segmentation of passengers, the airline would need to charge all passengers a different rate. This would be operationally infeasible. The case of first degree price discrimination is instructive, however, in that it represents an upper bound on the revenue achievable by an airline. Though not an attainable goal, this type of price discrimination provides a limiting case against which valuable comparisons can be made.

2.1.2 Second Degree Price Discrimination

Second degree price discrimination, also called self-selecting, seeks to identify the willingness to pay of passengers through their purchasing behavior. For example, those passengers having the greatest willingness to pay will demonstrate this by purchasing the least restricted fare product as a result of their preference for travel flexibility. Consumer self-selection is contingent upon a high degree of correlation between passenger sensitivity to fare product attributes (travel flexibility) and willingness to pay.

The array of fare products offered by airlines in each OD market can be seen as attempt to use second degree price discrimination to segment the demand population using their assumed sensitivity to travel flexibility. The most common devices used by airlines to achieve these segmentations are advance purchase, non-refundability, Saturday night stay, and round trip purchase itinerary requirements, in addition to availability limits.

Consider a simple example in which passengers must choose between the unrestricted product and a single discount fare product requiring a Saturday

night stay at the destination city. Passengers with the highest values of willingness to pay are assumed to prefer the unrestricted product because they are extremely sensitive to travel flexibility and unwilling to stay over the Saturday night. They are willing to pay a premium to ensure travel convenience. Conversely, passengers with values of willingness to pay lower than the unrestricted fare product price level (but above the discount fare product price level) are expected to be less sensitive to travel flexibility and thus, willing to stay over the Saturday night in order to secure the discount. If the self-selection scheme works, the two groups of consumers will identify their willingness to pay through their purchase decisions and be perfectly segmented.

The effectiveness of second degree price discrimination relies upon the ability of the airline to design and implement a differentiated fare product structure that capitalizes on the relative sensitivities of passengers and forces their segmentation by willingness to pay. Among the three types of price discrimination, the techniques of second degree dominate the industry. The basic fare product inferiority hierarchy that proliferates in most U.S. airline markets is the result of an attempt at self-selective passenger segmentation. Chapter One provided a detailed description of the types of fare products airlines use to segment passengers by their willingness to pay for air travel. All of these fare products can be modeled within the second degree price discrimination framework.

2.1.3 Third Degree Price Discrimination

Third degree price discrimination, also called index sorting, assumes the ability of the producer to segment the market *a priori* based on exogenous passenger characteristics. The exogenous sorting mechanism is expected to identify groups of the population who have similar levels of willingness to pay. Unfortunately, sorting passengers in this fashion is often difficult. Additionally, it may create two types of problems for the airline, lost revenue opportunities and revenue dilution.

Only a limited number of passengers within a given range of willingness to pay may be identified by the screening. Passengers not identified in the screening

process may result in lost revenue opportunities for the airline. On the other hand, the screening may falsely identify passengers who have values of willingness to pay significantly higher than expected. These passengers would have been willing to pay substantially more for air travel. As a result of meeting the screening criteria, however, they are able to purchase the discounted fare product thus diluting the revenue of the airline. An airline must be careful to avoid the potential pitfalls associated with third degree price discrimination when selecting screening devices.

Airlines have, in fact, made use of *a priori* screening techniques, as the following example illustrates. In the airline business, senior citizens and students are often identified as groups of potential passengers with low levels of willingness to pay. Thus, offering senior or student discounts allows the airline to serve these sections of the population at a price they are willing to pay without diluting the revenues paid by the non-senior/non-student segments of the population. The optimistic airline manager believes that the initial *a priori* segmentation of students and seniors results not only in identifying a group of passengers by their low values of willingness to pay but also in identifying a group possessing a higher degree of travel flexibility.

Not all students and senior citizens, however, have low values of willingness to pay. Depending upon their trip purpose and income levels, members of both the student and senior populations may possess some of the highest values of willingness to pay. Recall that Lee Iaccoca is a senior citizen. In addition, the populations of seniors and students represent only a fraction of the total population with relatively low values of willingness to pay. This example illustrates both the potential for revenue dilution and the lost revenue opportunities that often accompany the pure application of third degree price discrimination techniques to air travel populations.

To minimize revenue dilution and lost revenue opportunities, the promotions that are offered by the airline marketing department must be carefully designed. Their purpose is to stimulate travel from passengers with greater sensitivity to price level in order to increase load factors while not lowering yields as a result of dilution from passengers with higher levels of willingness to pay. Appealing to portions of the population who are willing to meet the inconvenient travel

51

restrictions in order to secure a lower price level who would not travel otherwise is the goal of airline marketing managers. The fear of revenue dilution from passengers of all kinds who have high levels of willingness to pay, however, is the primary one of these managers and thus, drives the fare product designs of airline marketers. Accordingly, it is common for airline fare product promotions to excessively limit the screening device and lose some revenue potential rather than risk diluting revenues.

Using the techniques of segmenting the market with an *a priori* signal about willingness to pay does not prevent the airline from applying self-selection techniques as well. The exogenous screening techniques used by airline marketers are invariably accompanied by restrictions on the flexibility of the passenger itineraries in accordance with the preferred tradeoff of airline management. The self-selection devices are applied to the discount fare products offered to the exogenously segmented populations using different product restrictions such as specific flight availability, booking limits, service guarantees, purchase restrictions, or combinations of these techniques. In other words, more heavily restricted fare products are offered to the segmented populations.

2.1.4 Second and Third Degree Price Discrimination Techniques in Practice

The combination of second and third degree price discrimination techniques is commonplace in the airline industry. In fact, the application of sophisticated yield management techniques his virtually eliminated all forms of pure third degree price discrimination. A closer look at the some of the latest available airline promotions provides insight into the forms of price discrimination practiced that have previously been assumed to be third degree price discrimination but, in fact, are combinations of second and third degree price discrimination. Student discounts are the first topic addressed.

At present, the most widespread student discounts offered are on Continental Airlines in conjunction with American Express. A similar offer between Chase Visa and USAir also exists. All who qualify for the discounted travel are students who hold American Express charge cards. This attempt at exogenously identifying a percentage of the population expected to have low values of

willingness to pay for air travel and high travel flexibility offers a clear example of one current application of third degree price discrimination.

In addition to the initial screening, however, passengers using the discounts are required to purchase a round trip itinerary, to stay at their destination over a Saturday night and to complete the trip within a seven day/six night period. The credit card-holding students are also subject to availability limits identical to those imposed on the lowest priced leisure travelers. The value of the students' travel itineraries may be reduced since the most convenient flights may not be available to them.

Another common third degree price discrimination technique is segmenting passengers traveling with children, most recently employed by Northwest Airlines last summer prior to the large scale fare war. In the case of "kids fly free" promotions, the accompanying child is subject to purchase restrictions and availability constraints based on the adult traveler's itinerary and fare product choice. Thus, in order to qualify for the free ticket, the child may be required to stay over a Saturday night, purchase a round trip ticket, and possibly be denied the first choice of travel itinerary because of availability limits depending upon the fare product purchased by the adult. In this respect, both adult and child exhibit the self-selection behavior intended by the airline as indicated by the fare product that they purchase.

The example closest to pure third degree price discrimination pertains to price reductions for young people. People under 25 years of age are offered a discount on the Delta and USAir Northeast (BOS-NYC-WAS) shuttles with no apparent restrictions other than proof of age. Differentiated fare products are generally not offered for sale on the Northeast shuttle routes. One fare product with one price level faces all. Differentiated fare products are not introduced for the youth travel population since they do not exist for the population-at-large. Even availability constraints do not readily apply to the flights due to the extreme overcapacity serving these markets. Outwardly, this scenario would seem to describe pure third degree price discrimination. Upon closer inspection, however, the youth fares are only available on the off-peak flight departures causing young people desiring the lower fares to self-select in terms of flight availability.

In practice, self-selection is never removed from the process of price discrimination in airline markets. Even the passengers initially screened using third-degree price discrimination techniques are subject to some self-selection criteria. Overall, in the airline industry, the scope of index sorting is quite limited. Promotions such as "kids fly free", student discounts, and other attempts to identify portions of the population with low values of willingness to pay affect only a small segment of the population when compared to the techniques of identifying passengers exclusively by their sensitivities to fare product attributes. Pure self-selection is, by far, the most widespread method of identifying and segmenting the consumer populations.

2.2 Previous Justifications of Price Discrimination in Airline Markets

On the surface, it is not clear to what extent the price discrimination that airlines practice is justified. The existence of price discrimination has direct implications on efficiency. Questions of equity, however, are not so obvious. Closer scrutiny of the factors underlying both airline supply and demand at the OD market level will provide some insight into the equity of airline price discrimination. The previous work focusing on the justifiability of airline price discrimination is explored in this section. A discussion of the economic efficiency ramifications of price discrimination is postponed until the following section.

The first justification of airline price discrimination is contingent upon the existence of scale economies in the industry. In a decreasing cost industry, price discrimination may be necessary in order to meet total costs. There may be no single price at which an airline can operate at a profit. In this case, a system of price discrimination may in fact be an equitable method for the product to be made available on the market while allowing the producer to meet costs. The existing airline fare product structure and the price discrimination that prevails simply may be a response to the presence of decreasing costs in the industry.

If airlines are not a decreasing cost industry, however, price discrimination may still provide some benefit to society. A higher level of service may be available to passengers from the practice of price discrimination. The increased frequency that may result from the addition of capacity enabled by the contributions to

54

overhead made by discount passengers benefits the passengers paying higher prices if they are sensitive to service frequency, as has been hypothesized. There are, in fact, potentially positive impacts of airline price discrimination even in the absence of scale economies.

Belobaba (1987) argues that the differential fare levels that are offered by airlines can, in fact, result in a Pareto optimal situation in which no consumers are worse off than they would have been in the absence of differential pricing. A detailed discussion of the existence or absence of Pareto optimality has not been presented in the literature. The generalized cost model developed in this research addresses the question of Pareto optimality in airline fare product differentiation within the current airline industry environment. Later in the dissertation, the conditions required for Pareto optimality are shown to be quite restrictive and unlikely to exist in the market.

Although Pareto optimality has not been fully addressed in the existing literature, several studies have shown reductions in average industry fare levels during deregulation. Schwieterman (1985) demonstrated that the price levels for similar fare products decreased during the early years of U.S. airline deregulation. The nature of the deregulated environment has changed since then, however. The GAO (1990) has demonstrated that average fare levels have decreased since deregulation although they have increased since 1985. Morrison and Winston (1990) have predicted that the fare levels paid by travelers in the deregulated environment are, on average, below what they might have been had the CAB fare formula been applied in the current environment. Some of these benefits may be attributable to other results of deregulation, however, and not exclusively the ability of airlines to product differentiate.

Frank (1983) contends that price-insensitive buyers stand in the way of the consolidation of airline flights into more economical (from a cost point of view) services. His hypothesis is also based upon the assumption that economies of scale exist in airline operations. Among the examples cited by Frank were the benefits of cost efficiencies resulting from demand consolidation. For instance, it is less expensive to operate a large aircraft fewer times than a small aircraft more frequently when providing an identical number of seats in an OD market. On a per passenger basis, such demand consolidation clearly results in operational

55

cost efficiencies. Intuition such as that offered by the example of demand consolidation strongly suggests that scale economies exist in the provision of airline services.

Contrary to the claims of Frank and intuition, several empirical studies have concluded that there are no economies of scale in providing airline service (Keeler, 1972, Douglas and Miller, 1974, and Christensen, Caves, and Tretheway, 1984). These models use standard unit-times-distance (UTD) measures, such as passenger-miles, to characterize transportation output. UTD output models implicitly assume proportional traffic flows across the network (Winston, 1985). This proportionality is highly questionable in airline markets as a result of the superior quality of service that is offered by non-stop carriers. Flight legs routinely carry a disproportionate number of non-stop passengers compared to connecting passengers. The conclusions made in these studies should be viewed skeptically since all have based their analyses upon the assumption that a generic measure of output correctly characterizes demand levels. The questionable proportionality assumption coupled with the counterintuitive findings make it difficult to place confidence in the studies showing no economies of scale.

Multi-product output models which model transportation product as flows over specific OD markets would provide a more accurate representation of reality. The definition of transportation product as a vector of flows offers a more intuitive characterization. Such models have been successfully applied to motor carriers and rail (Wang and Friedlaender, 1981, and Jara-Diaz and Winston, 1981). Unfortunately, no application of multi-product output models for airlines has been performed. Although it is beyond the scope of this research, it is hoped that the techniques of multi-product output can be used to arrive at a more robust measure of economies of scale in scheduled airline service. Perhaps then, a believable assessment of scale economies existing in airline markets can be made.

If economies of scale do, in fact, exist in the industry, it benefits airlines to consolidate their operations and offer fewer flights of larger aircraft from a cost standpoint. Price-insensitive consumers, however, are inflexible when faced with frequency reductions. According to Frank :

Certain buyers have strong preferences for special features and will readily abandon sellers that do not offer such features in favor of those who do, even if it means paying higher prices.

In this context, frequency of service is the special feature for which the price-insensitive passengers are prepared to abandon specific airlines (sellers). The aversion of price-insensitive travelers to frequency reductions forces the airline to offer, at a minimum, daily service in markets in which they hope to effectively compete. Daily, or even twice daily, service frequency is widely thought to be the absolute minimum requirement for competition in most U.S. domestic markets. Depending on the market, the frequency requirement may be even greater. Passengers that are sensitive to frequency of service impede the ability of airlines to consolidate their operations and thus, their ability to exploit the economies of scale available to them through flying larger aircraft less often. This renders the scale economies of demand consolidation unachievable, adding unnecessarily to system-wide costs.

Frank contends that differential pricing is not the result of market imperfections as has been stated in the economics literature but rather, a response to the economies of scale in OD market service that are foregone in order to provide the level of service that is desired, if not demanded, by price-insensitive consumers. Thus, it is plausible to view airline differential pricing as an attempt to distribute the burden of overhead recovery more equitably among different passengers. In this light, airline fare product differentiation may be viewed as a device forcing passengers to pay for the overall costs that they impose on the system.

The requirements of the passengers who travel most frequently and value the frequency of service the most dictate the long run aircraft procurement decisions. Thus, the price-insensitive passengers who make frequent service a requirement to garner their business drive up fixed airline costs in the long run. More efficient fleet acquisitions are foregone to provide service at the desired frequency of the price-insensitive passengers. This would imply that price-insensitive passengers should, in fact, pay more because of the larger long run costs that they impose on the system.

The conclusion of Frank that inefficiencies exist in the market as a result of by-passed economies of scale rather than market imperfections have been shown to be incorrect in the economics literature (Borenstein, 1992). Frank claims that there exists no price discrimination since price-insensitive passengers drive up the fixed costs of the carrier by forcing them to forgo available scale economies. Simple microeconomics tells us that long term costs do not figure into the marginal cost calculation. Since the strict definition of price discrimination is differential markups over marginal costs, the fixed costs of the airline have no effect on price discrimination which considers only price level and short run marginal cost.

For a fixed aircraft schedule, the differential markups over marginal cost facing consumers are clear and cannot contradict the findings of exchange inefficiencies by economists who have studied the problem (Borenstein, 1992). However, the societal impacts and the equitable distribution of total overhead cost distribution remain indeterminate. An improved model of cost distribution is required to correctly allocate the actual costs placed on the system by different passenger types. It is beyond the scope of this research to develop such a model, although it is clear that a tradeoff exists between the efficient level of production and the level of service desires of certain empowered consumer segments. Suffice it to say that these segments place upward pressure on costs and attempts to force them to pay a larger share of the overhead burden are not entirely misguided.

2.3 Efficiency, Yield Management, and Airline Fare Products

The question of efficiency in pricing in airline markets is now examined. It is the intention of this section to demonstrate that airline fare product differentiation and yield management techniques are not simply a necessary evil for airline survival but actually the tools for an efficient and desirable allocation of resources. First, the most common type of efficiency cited by economists, efficiency in exchange, is discussed as it relates to airlines. The ability of airlines to achieve exchange efficiency is then addressed. Next, the topic of efficiency in the allocation of a fixed provision of airline seats is presented. The combination of the existing differentiated fare product structure with airline yield management techniques provides a clearer look at the efficiency that currently

exists in airline markets. Finally, it is shown that allocative efficiency can only realistically be achieved when differential (and likely discriminatory) price levels exist in the market.

2.3.1 Exchange Efficiency

In neoclassical microeconomic terms, to obtain the best (most efficient) allocation of resources in the market, the production of goods must achieve exchange efficiency for all goods in all markets. If exchange efficiency were to exist then no one in the market could be made better off by producing or consuming any more or less of any particular good given a fixed amount of resources. Such a situation results in the optimal allocation of resources: using the least amount of resources to achieve the greatest amount of utility. In the case of airlines, attaining exchange efficiency is not, however, a feasible alternative as this section will demonstrate.

Setting P = SRMC

Economists routinely express concerns about departures from marginal cost pricing. This is because, in order to have a pricing system that is economically efficient in exchange, price level (P) must be set equal to the short run marginal costs (SRMC) of production for all goods and services in the market. Overall market exchange efficiency can be achieved only if all goods and services are sold at price levels equal to their marginal costs. Hence, existence of exchange efficiency is contingent upon airlines, as well as all other industries, practicing marginal cost pricing. This section describes why marginal cost pricing is not achievable in airline practice.

It is convenient to think of the SRMC faced by an airline in terms of the short run variable costs of carrying an additional passenger on the aircraft. Constant SRMC are often assumed to exist for airlines since the extra fuel, baggage handling, ticketing, and meal service that the airline must provide to serve each additional customer should be approximately equal on any given flight departure. With an aircraft type preassigned and no possibility of a change, the fixed marginal cost assumption for the flight departure is defensible, provided the flight only serves a single OD market. Unfortunately, the nature of airline

operations is far more complex in reality and the more complexity that is introduced, the more likely a departure from constant SRMC becomes.

The constant SRMC assumption is valid if a fixed and unalterable capacity exists in the isolated OD market. The ability of an airline to vary the capacity on a selected route makes the calculation of the marginal cost of carrying an incremental consumer far less straightforward. Short run aircraft capacity can be varied in two ways, either by flying an additional aircraft or by changing the size of the aircraft, both of which have impacts on the SRMC calculation.

First, the incremental cost of additional capacity resulting from flying an additional aircraft, or adding an "extra section", is considered. The incremental unit of production in the airline industry is the assignment of an additional aircraft to a route in order to carry more passengers. The incremental unit of sale, on the other hand, is an individual seat on an aircraft. With an aircraft of 100 seats, the incremental cost of the 100th passenger requesting service will be quite small, ignoring opportunity costs. Previous studies have estimated the incremental cost of serving an additional passenger provided that a seat is available at less than $25 (Air Transport World, 1986). The incremental cost of the 101st passenger, on the other hand, can be as much as the cost of flying the additional aircraft which, by any calculation, is quite large. This cost is likely to be greatly in excess of what any single passenger is willing to pay for service. Provided that the second aircraft is going to fly, however, the incremental cost of the 102nd passenger drops to near zero again (provided, of course, that the extra section is at least a two-seater).

The ability to switch aircraft further complicates the calculation of the marginal costs associated with a given flight departure. For instance, while the marginal cost of the 101st passenger is nominal on a 120 seat aircraft, it becomes quite large when the route is served by a 100 seat aircraft. The aircraft scheduling decision greatly impacts the marginal cost calculations for a given passenger. The opportunity costs associated with the allocation of aircraft capacity to a given route adds additional complexity to the marginal cost calculation. In the absence of the fabled "rubber aircraft" whose operating costs decrease proportionately with the number of passengers carried, marginal cost pricing is quite difficult when faced with the possibility of switching aircraft.

When expanding the marginal cost calculation for a single aircraft departure in an isolated OD market to multiple origin to destination markets, the calculation becomes even less clear. The hub and spoke network systems that pervade the U.S. domestic market require such considerations for realism. The marginal cost of providing service to a non-stop passenger, for instance, must incorporate the opportunity cost of not serving any of the potential passengers using the flight departure to connect to another flight.

While the marginal service provision for airlines appears to be the additional flight, the addition is somewhat deceptive. In reality, the true marginal service provision is the flight cycle. Flight cycles are required to provide a continuous, repeating schedule of airline services. Thus, an airline can increase its service provision by a minimum of two flights (the simplest cycle). In addition to network continuity considerations, the aircraft are required to cycle at some point in their round of flights for scheduled maintenance at a maintenance base.

It is clear from the above examples that the capacity provision decision in terms of both the size and number of aircraft in service strongly influences the marginal cost calculation. Network and multiple OD market service effects complicate the SRMC calculation further. Suffice it to say, the marginal cost calculation for any OD market is quite a complicated and unstable process. Kahn (1970) points out that, in reality, pricing at marginal cost may not be efficient if the calculation of marginal cost consumes more resources than could be saved by pricing at marginal cost. This is clearly the case for airlines.

Equilibration of Supply and Demand: Techniques for Narrowing the Gap

Another requirement of exchange efficiency is the equilibration of supply and demand. Supply-demand equilibration faces two seemingly insurmountable problems : the need for a perfectly competitive marketplace so that equilibrium can be reached and the ability to control incremental transportation output to perfectly match prevailing market demand levels. With U.S. industry-wide load factors at levels below sixty-five percent in 1989 and 1990 with no major carrier exceeding seventy percent in either year, achieving equilibrium appears to be unattainable for airlines in practice (Aviation Daily, 1991). Moreover, the limited number of carriers that operate in most domestic markets makes achieving

perfect competition unlikely. Some techniques which attempt to narrow the gap between air transportation supply and demand are available. The usefulness of these techniques is widely varied, however, and subject to numerous practical constraints when implemented. This section gives a detailed discussion of the methods used which attempt to reconcile supply and demand.

Since the equilibration of supply and demand in air transportation is very rare because of the marginal service provision requirements detailed above, airlines are put in the unenviable position of stimulating or turning away potential business almost continuously. Airlines employ various techniques to narrow the gap between supply and demand. On the supply side, improvements to the flight schedule through departure changes and aircraft switching may help to better match supply and demand on flights by increasing traveler convenience. Turning to the demand side, pricing policy can be used to manage demand given a fixed supply. Higher prices may serve to direct demand away from the most popular flights. These techniques may be attempted before the booking process begins for a given flight which is beneficial since practical constraints prevent any large scale changes in either the flight schedule or the published price levels close to flight departure.

Offering departure times that are convenient for more travelers may help to stimulate demand and generate incremental revenues for the carrier. Flights that are shown to be among the least popular may be switched to more convenient times by the scheduling department. The scheduling flexibility of an airline is limited in the short term, however, for many reasons. Airlines must attempt to adhere to their published schedules to allow their passengers adequate time to plan their travel. Reality dictates that airlines must operate with a flight schedule that is fixed in the short run, although it may change approximately monthly. The available seats are sold to passengers based upon the fixed capacity provision in each OD market as published in the current flight schedule. These considerations cannot be ignored when attempting to equilibrate supply and demand using aircraft schedules.

In addition, the generation of convenient departure times to stimulate demand is complicated by the physical constraints of aircraft operations. While consumers may unambiguously prefer 9:00 AM and 5:00 PM departures, not all aircraft may

fly at these times. Departure and landing capacity at peak times often impedes the most convenient departure times for all flights. As a result, only certain flights may receive peak departure and arrival times. Passengers have also been shown to prefer Friday and Sunday departures to Wednesday and Tuesday. Given a fixed aircraft allocation, however, it is quite expensive to maintain many extra aircraft to fly on these specific days.

Finally, airlines are unable to pick and choose the best departure times without considering the network effects. For instance, if a 9:00 AM EST departure from Boston to Detroit departs, an aircraft becomes available for service in Detroit at approximately noon. The cost of flying another aircraft empty to another location, or "deadheading", is generally quite large and not frequently done. In order for an aircraft to make sufficient contributions to justify its acquisition, it must maintaining a reasonable utilization rate. In this example, the aircraft should be put into use at Detroit sometime in the early afternoon, even though it may not be the most desirable departure time.

A system of peak-utilization pricing may be employed in order to direct price-sensitive passengers from high demand flights to less popular ones. The ability to equilibrate supply and demand using pricing, however, is limited in practice. Operational rules for the application of peak-utilization pricing are constrained by the ability of the airline to identify the high demand flights and apply higher fare levels to them. An airline might prefer to set a different fare for each flight in order to best direct (equilibrate) demand in this fashion. However, added complexity entailed may prove too great for the customers and ticketing agents alike.

In deference to practical constraints, simple day of week or time of day rules are often applied to peak/off-peak price differentials. This scheme may be counter-productive in certain markets which have different peaking behaviors. Careful attention must be paid to market-specific peaking characteristics. Since a plethora of fares are available to consumers which already include crude attempts at peak/off-peak pricing, a more complicated fare structure is probably not a feasible method for managing demand further. "Night coach" fares and other off-peak demand stimulating fares have been used in the past to direct passengers with low values of willingness to pay off of high demand flights and

onto lower demand flights. Even flight specific fares have been known to exist in airline markets. Lower flight specific fares are generally only offered at the most inconvenient times thus providing limited benefit for passengers competing for space on the most popular departures.

Even the most sophisticated attempts at demand direction are doomed to fall short of matching supply with demand. After the flight schedule and price levels have been fixed, popular flights will still experience excess demand while unpopular ones will depart with empty seats. After all of the attempts at equilibration of supply and demand have been made, the available seats on excess demand flights must still be rationed according to some allocative scheme.

Unable to engage in marginal cost pricing or obtain a supply-demand equilibrium that is efficient in exchange, another more operational method must be found for setting airline price levels and allocating the available resource. Hopefully, one that possesses desirable efficiency properties. In the next section, the allocative scheme prevailing in the deregulated U.S. airline environment is shown to possess such qualities.

2.3.2 Allocative Efficiency: Perhaps What is Really Desired

In the absence of an allocative scheme that is efficient in exchange, an alternative must be found. One such possibility is allocating the fixed aircraft capacity based upon the value that individual passengers place upon travel. In that case, an appropriate goal is allocative efficiency which is defined as making sure that a scarce resource is provided to the members of the population who intrinsically value it most. Allocative efficiency can be viewed as making the best possible use of a resource.

As would be expected, available seats on the most demand intensive (or valuable) flights are scarce and must be rationed out to passengers according to some allocative scheme. When a flight is near capacity it is likely that the purchase requests of some passengers have been denied. The denied requests result in lost revenue for the airline and decreased convenience for the rejected passengers. Consequently, the scarce resource must be allocated according to some decision rule on flights with excess demand.

64

If a first come first serve (FCFS) allocation process is employed, the passengers who have obtained seats on the flight may not, in fact, have valued the seat most. In this way, the situation facing air transportation consumers differs from that facing the producers of other goods and services. Under perfect information, passengers valuing the service more would make bookings earlier relative to passengers valuing the service less. Unfortunately, the value of air transportation may not be known to the consumer until close to the flight departure, making an early booking impossible.

While advanced planning and the ability to book a seat early may be highly valued by some, many of the consumers can be unsure of their exact travel itineraries until close to departure and still value air travel highly. Thus, although a FCFS discipline might drive many passengers valuing service highly to book earlier, the scheduling and planning uncertainty of others makes the allocation of the limited number of airline seats on excess demand flights to those who value them most doubtful. As a result, the nature of passenger air transportation demand makes a rationing of airlines seats in a market employing a FCFS allocation unlikely to achieve allocative efficiency.

Employing the techniques of revenue management offers an alternative to a FCFS discipline. Revenue management uses differentiated fare products priced at varying levels and booking limits to provide a flexible, flight-specific method of seat inventory allocation. A beneficial result of the revenue management function is directing demand to lower load factor flights making the best use of available capacity. In this way, the allocative efficiency of the system can be preserved or enhanced. Attention must be paid, however, to the tradeoff between demand direction and passenger service rejections. This section highlights the methods currently used for achieving an allocatively efficient process. The use of price discrimination and revenue management techniques in this context are discussed.

Identifying Passengers by Willingness to Pay : Price Discrimination

When a resource is limited, a loss of efficiency would result if that resource were allocated to passengers placing a lower value on it than other passengers to whom it was denied. The best use of the resource would be to provide it to those

consumers who value it most highly. For airlines, heavily demanded aircraft seats should be allocated to the passengers who value the seats most highly (up to capacity). Unfortunately for the airline and allocative efficiency, passengers do not willingly identify themselves by the value they place on securing a seat for a given flight departure. Moreover, consumers do not identify their willingness to pay through the arrival process because of the uncertainty associated with travel itineraries (in contrast to most other goods and services). Airlines must thus employ screening procedures to identify passengers according to their willingness to pay. This enables the airlines to most efficiently allocate their resources given a fixed flight schedule. The techniques of price discrimination can be employed to identify passengers in the desired manner.

To enable an efficient allocation of the scarce resource, the important function of price discrimination is the identification of passengers by their willingness to pay. The limiting case of first degree price discrimination, in which each passenger is charged an amount equal to his willingness to pay, would guarantee an efficient resource allocation since the identification occurs at the unit of sale in the airline industry, the incremental seat. In this case, the airline would accept only the number of passengers equal to the capacity of the aircraft having the highest values of willingness to pay (provided that demand exceeded supply). For a static application, the airline would be able to choose the people with the greatest value of willingness to pay and let them on the aircraft while turning the others away. Of course, first degree price discrimination is unachievable in practice since, even if passengers could be exactly identified by their willingness to pay, people must know in advance of departure whether or not to expect a seat on the aircraft.

In practice, the identification of passengers by the value that they place on a specific air transportation service must be achieved by using an effective screening procedure which groups passengers by their willingness to pay. Both second and third degree price discrimination help the airline to identify the passengers who value the product most. Techniques of pure third degree price discrimination, such as identifying members of the population through their membership in a "gold" or "preferred" frequent flyer club, are quite fallible since consumers have different values for flights at different times. For example, it is not guaranteed that the passenger with the highest value of willingness to pay on

this week's Friday PM departure will not be among the passengers with the lowest values for next Friday's PM departure. Similarly, business travelers among the most price-insensitive when the company is paying may become highly price-sensitive when traveling on leisure trips for which they are paying. The limited scope of third degree price discrimination techniques limits their usefulness in creating an efficient allocation.

As previously motivated, the screening procedure that has been found most effective in airline practice is that of second degree price discrimination, or self-selection. Self-selection techniques are the most far reaching and effective method to achieve a passenger value identification. The array of differentiated fare products offered to consumers allows passengers to identify their value of willingness to pay for every airline ticket purchase. Moreover, the fare product structure allows passengers to secure their desired priority level each time that they purchase air transportation fare products. The self-selection enables a resource allocation based on the current trip requirements of the individual, thus, promoting efficiency through an accurate value identification screening.

Fare Product Structure Contribution to the Efficient Allocation : A Closer Look

The underlying structure of fare products required to facilitate an efficient allocation of the limited resource of available seats is now addressed. The required fare structure for allocative efficiency is one that identifies and segments passengers by their willingness to pay thus allowing the airline to fill the capacity with the passengers having the greatest intrinsic willingness to pay. Issues of price levels and purchase restrictions must be addressed to identify a fare product structure conducive to allocative efficiency. First, the properties beneficial to effective second degree price discrimination are outlined. The potential pitfalls of the self-selection techniques are then addressed. Finally, the impacts on allocative efficiency of identification and segmentation policies (and failures) are highlighted.

The techniques of self-selection attempt to stimulate demand from passengers with low values of willingness to pay while not diluting the revenues received by passengers valuing the product more. To stimulate traffic using the techniques of second degree price discrimination, the cost incurred by a passenger

associated with accepting the purchase restrictions plus the price level must not exceed the intrinsic value of travel on that flight. On the lower demand flights, the costs to passengers targeted to a self-selection stratum should be as low as possible to stimulate as much demand as possible. The ability to stimulate incremental traffic is less important on higher demand flights. To stimulate the most traffic, the costs associated with the imposed restrictions should be minimized for the target population.

To prevent revenue dilution, the purchase restrictions on the lower priced fare products must discourage passengers with higher values of willingness to pay from "buying down". The fare differentials should be kept low to limit the incentive for passengers to divert to lower-priced fare products. In other words, the cost associated with additional restrictions for the passengers targeted to any stratum must exceed the price differential between the target fare product and all more-restricted fare products. If this is not the case, passengers will purchase one of the lower-priced fare products. To minimize revenue dilution, the costs associated with accepting the purchase restrictions on a discount fare product should be very high for passengers willing to purchase higher-priced fare products in the absence of discounts.

When setting fares, the price levels must be strictly decreasing with increasing purchase restrictions. Any departures from pricing the fare products in this fashion would result in a case of simple domination. Price levels allow passengers to self-select on the basis of their sensitivity to fare level. In this way, price level can be thought of in the same way as other purchase restrictions. Attention must be paid, however, not to drop the price low enough to make the fare product attractive to portions of the population valuing travel significantly more.

The restrictions associated with lower-priced fare products should be designed to prevent passengers with higher than targeted values of willingness to pay from diverting down while imposing the smallest possible cost on passengers in the target range. Since it is commonly assumed that passengers with greater values of willingness to pay also have greater sensitivity to travel flexibility, most discount fare products currently available on the market strictly limit travel itinerary flexibility in their restrictions. A judicious selection of purchase

restrictions along with the correct price levels helps to prevent improper self-selection and revenue dilution for the airline. A balance between highly effective/low cost fences, however, may be difficult to achieve. As a result, airline managers may be required to make tradeoffs between discouraging demand and diluting revenue. The need for effective self-selection techniques and price levels is apparent.

Fencing techniques are quite fallible due to demand variability and changing passenger characteristics. When the sorting techniques used by second degree price discrimination are not perfectly effective, departures from allocative efficiency may occur. Discouraged demand may result when the self-selection criteria used imposes too great a cost on passengers who ordinarily would choose to travel using one of the less restricted fare products. For instance, if a passenger was willing to pay $750 provided that he was not required to stay over a Saturday night but the lowest published non-excursion fare was $800, he may choose not to travel even though there were many excursion fare products available for significantly less than $750. The lack of acceptable fare products in the price range of passengers like the one above would result in an inefficient allocation of the available resource from the failure of self-selection.

If passengers beat restrictions and divert to lower-priced fare products, allocative efficiency is unaffected. Not surprisingly, the ability to divert to lower-priced fare products can only benefit the passengers having the highest values of willingness to pay. Even if certain price-insensitive consumers can beat restrictions and purchase lower-priced fare products, the current seat allocation discipline still guarantees them a seat if there is one available provided that they have the highest value for it among the remaining unserved passengers. Thus, allocative efficiency is upheld even with passenger diversion to lower-priced fare products.

The self-selection techniques must be equated with the passenger demands to preserve the scarce capacity for consumers who value it most. The fare product prices, in turn, cannot be so high as to discourage passengers who would have traveled and contributed to fixed costs had prices been lower. The demand levels for individual flights often vary by season or even by week necessitating a

flexible fare product structure. Thus, airlines seeking to efficiently allocate scarce seats must match fare levels with travel values to the best of their ability.

Offering a wide range of fare products with different price levels and appropriate purchase restrictions would help to account for flight and seasonality variations. Such a fare structure, if achievable, can better simulate first degree price discrimination. A fare product structure that is too complex, however, may lead to consumer or travel agent confusion. Practical constraints also require that the same basic fare products be offered on all flights to limit consumer confusion. Thus, what is theoretically best for allocative efficiency may not be completely feasible in practice.

Because the airlines have nearly complete control over their seat inventory, offering a wide range of fare products is desirable since sales can be controlled using booking limits. A proper application of yield management techniques allows the airline to protect seats for higher valued passengers when demand is high and to not discourage (or prevent) lower valued passengers from filling less demanded departures. The next section demonstrates how booking limits help the airline to preserve an efficient allocation when available seats are scarce.

2.3.3 Using Yield Management Techniques to Achieve Allocative Efficiency

Self-selection alone would approximate an efficient allocation for the static case. Airlines, however, face a dynamic booking process in the presence of which they need to protect scarce inventory for passengers with higher values of willingness to pay. Also, the nature of the self-selection devices employed by airlines, in particular advance purchase requirements, require a dynamic seat allocation. The practical constraints of the dynamic booking process require the seat inventory control capabilities offered by computer reservations systems to be used to preserve an efficient resource allocation. With proper reservations control, the application of nested booking limits ensures that increasingly less restricted fare products have increasingly greater availability reinforcing the airline fare product inferiority hierarchy.

A perfect (hindsight) application of yield management techniques to a nested seat allocation approximates an efficient allocation of the available seats on full,

70

or near full, flights. In the limiting case of first degree price discrimination, the efficient allocation is guaranteed with optimal booking limits. Recalling the yield management techniques described in Chapter One, the important elements of controlling the airline seat allocation on a given flight departure are the booking limits set for the individual fare products and the overbooking factors set to prevent seat spoilage from passengers who do not show up for the flight departure.

Booking Limits

Self-selection helps to identify strata of passengers with like values of willingness to pay for a seat on the aircraft. Instead of being able to identify every passenger by the value that they place on travel, the airline assumes that passengers are grouped into ranges of willingness to pay. The identification is limited, however, by the effectiveness of the self-selection devices. Airlines may only have a vague idea about the size of the strata into which passengers are grouped. Thus, proper demand forecasting and booking limits are required due to the high variability of passenger demand and the imprecision of self-selection techniques.

Booking limits take on great importance in the efficient allocation of aircraft seats when demand exceeds supply. If passengers are segmented into mutually exclusive groups by willingness to pay, the resource allocation is efficient if the price levels and booking limits are set correctly. Unfortunately, departures from efficiency may happen due to the imprecision associated with price discrimination and the inability to set prices that will equilibrate supply and demand.

Assume for the sake of exposition that an airline has determined the demand for travel on a given flight departure as exactly 150 passengers. The passengers have been perfectly identified in three fare product strata each having 50 members. The members of the fare product 1 stratum are assumed to have the highest values of willingness to pay while those in the fare product 3 stratum have the lowest. The passengers within the same fare product stratum have values of willingness to pay in the same range although they are not identical. Price levels are strictly decreasing from fare product 1 to fare product 3.

For a 100 seat aircraft, if no booking limits were applied, there would be a risk that the aircraft would not serve the passengers valuing travel most highly. The optimal allocation of seats to fare products would be 50 to fare product 1 and 50 to fare product 2 with no fare product 3 passengers being served. To prevent the possibility of an inefficient allocation, booking limits can be applied to the flight. If the airline were to apply a booking limit of zero to the fare product 3 population, for example, the efficient allocation would be preserved regardless of the arrival process. This simple example illustrates that the application of booking limits to the fare products is required to ensure allocative efficiency for the aircraft on excess demand flights.

Any aircraft having over 150 seats would provide excess supply given the demand levels in the example. No booking limits are required to obtain an efficient seat allocation since all passengers are served on excess supply flights. For a 150 seat aircraft, the price levels for the differentiated fare products have been set optimally for the single flight example. As a result, supply and demand have been equilibrated and there is no need to apply booking limits to the aircraft seat allocation. The case of supply/demand equilibrium is the exception rather than the rule in the airline industry which is characterized by heavy peaking and seasonality.

When the route is served by a 120 seat aircraft, allocative efficiency is no longer guaranteed even with the application of optimal booking limits because of the imprecision within the fare product 3 stratum. Under optimal booking limits, 50 fare product 1, 50 fare product 2, and 20 fare product 3 passengers are taken and revenue is maximized for the airline. Allocative efficiency, however, is contingent upon the arrival of the first 20 fare product 3 passengers being those with the greatest levels of willingness to pay among all fare product 3 passengers. Although possible, such an arrival process is unlikely. This demonstrates the departure from allocative efficiency that may occur in the flight resulting from the imprecision of self-selection techniques even under optimal booking limits. It is interesting to note that the departure from efficiency occurs only for the lowest-priced fare product having non-zero booking limits in the deterministic case and varies inversely with stratum size.

If the price level is set below a value which translates into demand exceeding capacity, then more passengers will arrive than can be accommodated. There is the potential for the loss of efficiency for the lowest priced fare product. In the deterministic case, it is clear that no loss of efficiency can occur for any fare product other than the lowest with optimal booking limits (greater than zero). The loss of efficiency, however, can be mitigated as the lowest fare product stratum size decreases.

Consider next a second flight in the market, having a demand level of 45 passengers split evenly between the same three fare products. Assume that both flights are served by 100 seat aircraft and the demand level for the first flight is the same as before. The demand levels for the flight departures appear in Table 2.1. It is assumed that all passengers initially requesting travel on the first flight prefer it to the second and vice versa. All passengers wishing to travel, however, prefer to do so on either flight over the option of not traveling. After optimal booking limits have been set for the first flight, 50 fare product 3 passengers will have been denied seats. The fare product 3 passengers from the first flight will secure passage, however, on the second flight since there are only 95 passenger requests for the 100 seat aircraft. The passengers on the first flight will be the 100 with the highest levels of willingness to pay for that flight. Similarly, the 95 passengers on the second flight will have the highest values for that flight. This example illustrates the efficiency associated with demand direction using yield management techniques.

Flight #1				
	Fare Product 1	Fare Product 2	Fare Product 3	Total
Demand	50	50	50	150
Bkg Limits	100	50	0	
On Board	50	50	0	100
Flight #2				
Demand	15	15	15	45
Flt 1 Demand	0	0	50	50
Bkg Limits	100	85	70	
On Board	15	15	65	95

Table 2.1

Seats may be reserved for the highest value passengers on flights that are expected to be at, or near, capacity by applying booking limits to the discount fare product strata. The use of booking limits on the discount fare products allows passengers with higher values of willingness to pay to have a better chance of securing a seat on the highest demand flights. The limits prevent early-arriving passengers with low values of willingness to pay from displacing passengers valuing the seats more who arrive later in the booking process. The availability limits also minimize revenue dilution for the airline, in addition to increasing seat availability for price-insensitive passengers.

The application of booking limits to the discount fare products allows the airline to sell fewer discount seats on the higher demand flights while offering virtually unlimited discount seats on the most unpopular flights. With a correct application of airline booking limits, the passengers valuing the product less will be redirected to other lower demand flights which are below capacity and on which they are likely to be among the passengers valuing the service most. This provides for an efficient allocation of the resource on both types of flights. Thus, under the same fare product price and restriction structure, the available seats can be efficiently allocated on both full and empty flights. An imperfect application of booking limits might result in a denial of service to the high fare/value passengers since not enough seats were protected for them.

To this point, the self-selection techniques used in the example have been assumed to identify and segment the strata without error. As previously stated, when passengers divert to lower-priced fare products than expected and no booking limits are applied, an efficient allocation can be achieved. The application of booking limits, however, can impede an efficient allocation of the available resource. Passengers who self-select at a lower level of willingness to pay than targeted may displace passengers who would only be willing to purchase at the lower level, in the presence of booking limits. The application of booking limits to the flights may introduce negative effects on efficiency by discouraging some passengers who are extremely sensitive to travel flexibility if the seats targeted to them have already been purchased by passengers with higher values than targeted. This becomes increasingly problematic on higher demand flights with the application of booking limits to the self-selected strata if passengers with higher values of willingness to pay arrive first.

Overbooking

The airline refundability policy increases the variance associated with demand and increases the probability of seats flying empty or "spoilage". Seat spoilage is primarily the product of the airline allowing refunds to passengers who neither show up for the flight nor cancel their reservations (although unforeseen no-shows also affect the spoilage). In this way, the absence of a refundability penalty for the less-restricted fare products may actually reduce the efficiency of the seat allocation. The problem of "no-shows," as they are called, is best addressed using overbooking techniques.

When setting the overbooking factors for a given flight departure, the airline is trading off the costs associated with a denied boarding with those incurred by allowing a seat to fly empty. A seat flying empty when a request for service has been denied is clearly not an efficient resource allocation provided that there is an additional passenger willing to pay at least the SRMC of providing service to him. A denied boarding presents less of a problem for allocative efficiency than for the passenger. From a practical standpoint, however, an airline must be concerned about the effect of denied boardings in terms of goodwill and required compensation. In other words, the airline can use overbooking factors to maximize the probability of an efficient allocation by minimizing spoilage. Operationally, however, the airline must realize that more aggressive overbooking policies may lead to increased denied boardings and, potentially, passenger complaints.

Bumping Procedure : Efficient Auction

It is clear that the yield management techniques of overbooking and booking limits and the pricing functions of setting fare levels and fare product design are of vital importance to the efficient allocation of airline seats when they are scarce. Through an integrated effort of the revenue management function, an efficient allocation can be achieved with high probability. Departures from efficiency are inevitable, however, because of the high level of variability occurring in passenger demand. The corrective technique of auctioning off the seats on an aircraft can help to achieve an efficient seat allocation when the number of seats sold exceeds capacity (as the result of an incorrect overbooking factor

application). In this way, gate agents can contribute to the revenue management process and correct any inefficiencies introduced by the revenue management function.

Flights having sold too many seats, or "oversold" flights, require certain passengers not to board the aircraft since all passengers must be seated by law. U.S. law also requires airlines to offer compensation to passengers who are deplaned, whether voluntarily or involuntarily. In terms of efficiency, it is preferable for the passengers who surrender their seats first to be those with the lowest values of willingness to pay. To better identify passengers by their willingness to pay, airlines have used a system of volunteering in which certain passengers are offered the opportunity to surrender their seats to passengers without seat assignments.

In exchange for the use of their seat, the volunteers forgo travel on the oversold flight and are offered compensation. In lieu of offering cash, airlines most often offer free travel as an incentive. Since compensation is required for bumping, minimizing the cost of the compensation is of interest to the airline. If no passengers volunteer, the reward level is increased up to a certain level as in an auction. As departure time approaches and no volunteers have come forward when the maximum reward has been offered, then passengers are involuntarily deplaned so that the aircraft can depart.

Compensatory travel is restricted using some of the self-selection techniques described previously as well as being availability limited. The high level of restrictions placed on the "free" tickets used in the auction prevent their substitution with unrestricted fare products. This helps to mitigate the revenue dilution effects of offering the compensation. The level of priority given to the tickets offered in the auction is at the level of the frequent flyer reward fare products (possibly lower than the lowest booking class) or they can be used on a space available basis 24 hours before departure. Again, the revenue management system ensures the self-selection of the passengers regarding their choice to use the free ticket. It is the hope of the airline that the compensatory ticket encourages passengers to take a trip that they would not normally have made, or to travel with a paying companion and actually generate revenue for the airline.

Under certain conditions, the efficient auction of tickets to remove passengers from oversold flights can actually enable the airline to achieve revenues in excess of those obtained from an efficient allocation of the aircraft capacity. This, of course, depends upon the assumptions surrounding the reward offered in the auction. If the opportunity costs are near zero and only surplus capacity is allocated to the tickets, the auction can be a win-win situation. The passengers denied boarding are switched to lower demand flight departures and the airline receives the revenue for their passage with little additional cost. The offer of compensation has even been known to engender goodwill among certain passengers like the author.

While yield management techniques such as overbooking and booking limits are often criticized in the popular press, they serve a valuable function in terms of preserving allocative efficiency. In fact, it is only the failure of yield management techniques (i.e., booking limits, overbooking) and fare product differentiation that impedes an allocatively efficient use of the scarce airline resource. Even the fallibility of the yield management can be mitigated using techniques such as the efficient bumping procedure. In this way, the application of airline revenue management techniques preserve the efficient allocation of the scarce resource.

2.3.4 Equity Considerations for Allocative Efficiency

It should be noted that the efficiencies achieved through the allocative procedure do not extend to the maximum value of consumer surplus among those traveling on the flight. The passengers aboard the aircraft will be those with the highest intrinsic value for being on the flight, or, those who possess the highest opportunity costs associated with not being on the flight. The efficient allocation only guarantees that the resource is provided to those passengers intrinsically valuing it most. Measures of consumer and societal welfare are not optimized in this allocation scheme. This shortcoming may be of concern to policy makers. There are other considerations about product differentiation and yield management techniques, however, which may increase equity.

Frank (1983) claims that the greater frequency provision is the reason that price-insensitive passengers are paying differential mark-ups over marginal cost. This is not entirely true. The increased frequency alone would not satisfy the price-

insensitive consumer if he were not able to secure a seat on the aircraft at his most preferred times. An uncontrolled seat allocation would likely result in leisure passengers securing passage on the most popular flight departures in a given market. This would be to the detriment of the price-insensitive passengers who would not be allocated the resource that they value so much more than other travelers.

The differential mark-up over marginal cost that is paid by the price-insensitive passenger can be thought of as the premium paid in order to ensure an efficient allocation of the scarce resource, if one is needed. Since the passengers who place the highest value on availability are often those with the least advanced travel notice, the efficient allocation of the resource is most important to them when faced with the alternative of a FCFS, or similar, allocative discipline. Price-insensitive passengers want discount passengers to contribute to fixed costs in order to keep fare levels lower but do not want to be displaced by the discount passengers when seat availability is scarce. For this, the price-insensitive passengers must pay a premium. The overall fares that they pay, however, may be reduced in the long run by the contributions made to fixed (and joint) costs paid by the discount passengers. Or, alternatively, their level of service may increase due to greater frequency of service offered.

2.3.5 Efficiency Summary

Exchange efficiency is not attainable in the deregulated U.S. passenger air travel market. Failing exchange efficiency, a system providing allocative efficiency given the fixed provision of output is likely the preferred recourse. The effect of yield management systems on the allocation of the fare products possesses efficiency properties. It helps to ensure that those passengers who value service the most are reserved space on the aircraft. The value of this efficient allocation is less persuasive when low demand flights are considered. As a result of the overcapacity that plagues the U.S. deregulated airline market, the benefit of this seat allocation may be questioned. However, due to the variability of demand that dominates the industry, and the extremely high value that price-insensitive passengers place on their travel itineraries, an efficient allocation clearly has value to them.

2.4 Monopolistic Competition Models of Price Discrimination

Models of third degree price discrimination are widespread in the literature both in monopoly and oligopoly markets. For example, in monopoly markets, Schmalensee (1981) has discussed the effects of third degree price discrimination on price differentials, welfare, and passenger demand. Borenstein (1985) and Holmes (1989) have independently addressed price discrimination in oligopolistic markets using spatial models of monopolistic competition based upon the framework developed by Salop (1979). Since airline markets are more oligopolistic in nature, the review of the literature presented here focuses on the oligopolistic models.

Holmes limits his study to an analytical discussion of duopoly for pure third degree price discrimination which is not appropriate in the context of airline price discrimination. Borenstein, on the other hand, allows for a general number of competitors (with a case study on up to four) and considers both third and second degree price discrimination using simulation methods. As motivated above, the case of second degree price discrimination is most appropriate to the airline case. Thus, the work of Borenstein holds the highest level of interest to this research.

Borenstein assumes the ability to identify passengers by their willingness to pay (reservation prices) and their brand preferences. The impacts of sorting by reservation prices and brand preferences are quantified using simulation methods under different market structures. The paper claims that sorting by brand preference will result in greater market price differentials than sorting according to willingness to pay in monopolistically competitive markets. This hypothesis makes a great deal of sense when thought of in the context of the hub and spoke networks prevailing in the deregulated U.S. airline industry.

The hub and spoke system will inevitably lead to markets in which some carriers offer non-stop service while others can only provide connecting service. In addition, the peaking behavior in airline markets shows certain departure times to be more attractive than others. The underlying correlation between willingness to pay for air travel and sensitivity to service attributes tells us that

the level of service differential between non-stop and connecting service heavily influences passenger air carrier choice.

Heterogeneous product offerings or brands are allowed in the model which can be thought of in terms of different flight departures, and frequencies, and can be used to model product loyalties. These factors are quite important when considering demand for air travel services. For an airline to affect their brand perception, however, it may involve large costs in terms of marketing, scheduling, aircraft type and capacity. The ability of an airline to improve its brand preference rating is quite a difficult task that may involve large capital outlays and also adversely affect the operations and brand perceptions of the carrier in other OD markets resulting from level of service changes with fixed resources.

If airlines are able to effectively sort passengers using their reservation prices, on the other hand, profits can be improved without the large costs associated with attracting passengers away from other airlines. Thus, in the view of the airline manager, the identification and sorting of passengers by their willingness to pay is a more cost-effective method to identify the overwhelming majority of passengers. Clearly, neither sorting ability is available to an airline without significant error. The model does, however, indicate the effect on the cited factors if the airline were, in fact, able to sort without error.

Borenstein discusses the concept of value differential vs. price differential in his model of monopolistic competition when sorting by self-selection is introduced. The cost associated with accepting an inferior fare product was written as a function of brand preference and passenger reservation price. According to the notation of the paper, the cost, K, associated with accepting an inferior product is written as $K = K(A,c)$ where A is the reservation price of the consumer, and c is a measure of brand preference. Passengers are assumed to choose the low price fare product when $K(A,c) < P_H - P_L$ where P_H and P_L are the price levels for the full and discount fare products, respectively. The full fare product is chosen by those consumers having $K(A,c) \geq P_H - P_L$ or a value differential greater than or equal to the cost differential. The inclusion of the cost of travel inconvenience is important to the correct modeling of the airline consumer decision process.

80

In the paper, the calculation of a profit maximizing sorting mechanism is ignored. The reason cited being that the number of ways that a producer can segment the market is limited. This leads to a binary grouping of passengers that was exogenously determined. Thus, only the two fare class case was explored. As a result, the model does not provide airlines with any directly operational suggestions for improved price setting practices since airlines offer far greater than two fare products in most OD markets.

Borenstein correctly indicates that the demand for the two fare products is not, in fact, independent of the price differential between them. The discount and full fare products are not assumed to be independent in his modeling framework in contrast to most of the work done in airline yield management modeling. The sorting of passengers done in the paper is based on parameters such as the cost associated with accepting a discount product and its associated inconvenience as well as upon knowing their reservation price levels and brand preferences. This accurately models the type of self-selective sorting done by airlines.

The results of the analysis provided information on the effects of different segmentation and sorting strategies on passenger demand, welfare, entry, and the size of the price differentials existing in the market. The paper provides information that is more useful toward formulating regulatory policy than helping airlines concerned with setting price levels in their OD markets. The results of the analysis do not reflect the structure of airline fare products and market segmentation techniques. Rather, they provide a simple description of the interrelationships of brand preference, reservation price, demand and industrial organization.

In the theoretical research done by Borenstein and Holmes, the price levels that prevail in the market are not set by an airline exogenously, but rather, are the result of a Bertrand equilibrium between the producers of the different brands. These equilibrium price levels may offer an indication of the nature of price levels that will result from a specific market organizational structure, but offer little aid in the exogenous setting of price levels practiced by the individual airlines. As previously motivated, equilibration of supply and demand is not in keeping with the realities of the airlines limiting further the applicability of the modeling framework.

Borenstein and Rose (1992) have addressed the issue of price dispersion as an indicator of price discrimination using a spatial model of monopolistic competition similar to that used by Borenstein (1985). The aim of the analysis was to describe the reasons behind the discrepancies between price levels paid by consumers on the same carriers in the same markets. The analysis used factors of service, frequency, and market dominance to explain the reasons behind the discrepancy in fare levels. Again, the structure of airline fare products and the setting of their price levels were not explicitly considered in the analysis. The study is of limited use to airline managers who are more interested in influencing factors under their own control.

Borenstein and Rose also conclude that the level of price dispersion increases as the competition increases in the OD market. Although many reasons were cited, increased dispersion with more competition may be explained best by the reaction of airline yield management systems to overcapacity. When the level of competition increases in a market, an increase in OD market capacity generally occurs with it. The yield management system, in general, responds to large amounts of overcapacity with increased seat availability in the discount classes. If increased discount seat availability does not increase load factors given the current fare product structure, the pricing department may drop the market fare levels, leading to even greater price dispersion in the market.

None of the attempts to model price discrimination have explicitly considered the structure of airline fare products offered in the OD market. Rather, they have focused on the departures from marginal cost pricing and whether or not they exist in the market and what is their basis. At the level of the individual airline, such studies of industry organization are of limited relevance. Airline managers are more interested in determining the effects of the decisions over which they exercise control. Previous research has offered some insights into the effects of setting profit maximizing price levels in OD markets for the individual fare products. This research is introduced in the next section along with a critique of the assumptions that underlie the models, which provides a motivation for the research undertaken in this dissertation.

2.5 Monopoly Models of Airline Fare Product Differentiation

Previous attempts to provide profit-maximizing price levels under fare product differentiation have been attempted in the literature using a model of third degree price discrimination (Kahn, 1970, Belobaba, 1987, and Cross, 1989). The model used to describe airline fare product differentiation in these models is based upon a single flight in a single OD market. The model appears graphically in Figure 2.1. The demand is static and comes from a single population. Thus, airline passengers are only concerned with obtaining a space on the aircraft as indicated by their willingness to pay for air travel.

Passenger willingness to pay is unambiguously described using the single demand function appearing in the figure. The short run marginal cost of serving an additional passenger has been assumed to be constant in the model. The average total cost curve is above the demand function at all points. Because of high average costs, the airline is unable to cover costs at any single price level. Thus, in order to offer the flight over the long term, the airline is forced to price discriminate.

The airline is assumed to possess the ability to price discriminate using an infallible screening technique. Using the sorting mechanism, passengers are grouped by their willingness to pay into mutually exclusive groups. It is assumed that passengers grouped into a stratum are prevented from purchasing any other fare products because they are subject to a perfect fencing mechanism. Additionally, the strata are decided by the setting of the price levels for the market. This assumption about the ability of the airline to identify and segment passengers is greatly unrealistic.

The model presented in the example is a third degree price discrimination model for a monopoly which relies upon the ability of the airline to perfectly identify

83

Monopoly Price Discrimination
Three Fare Product Example
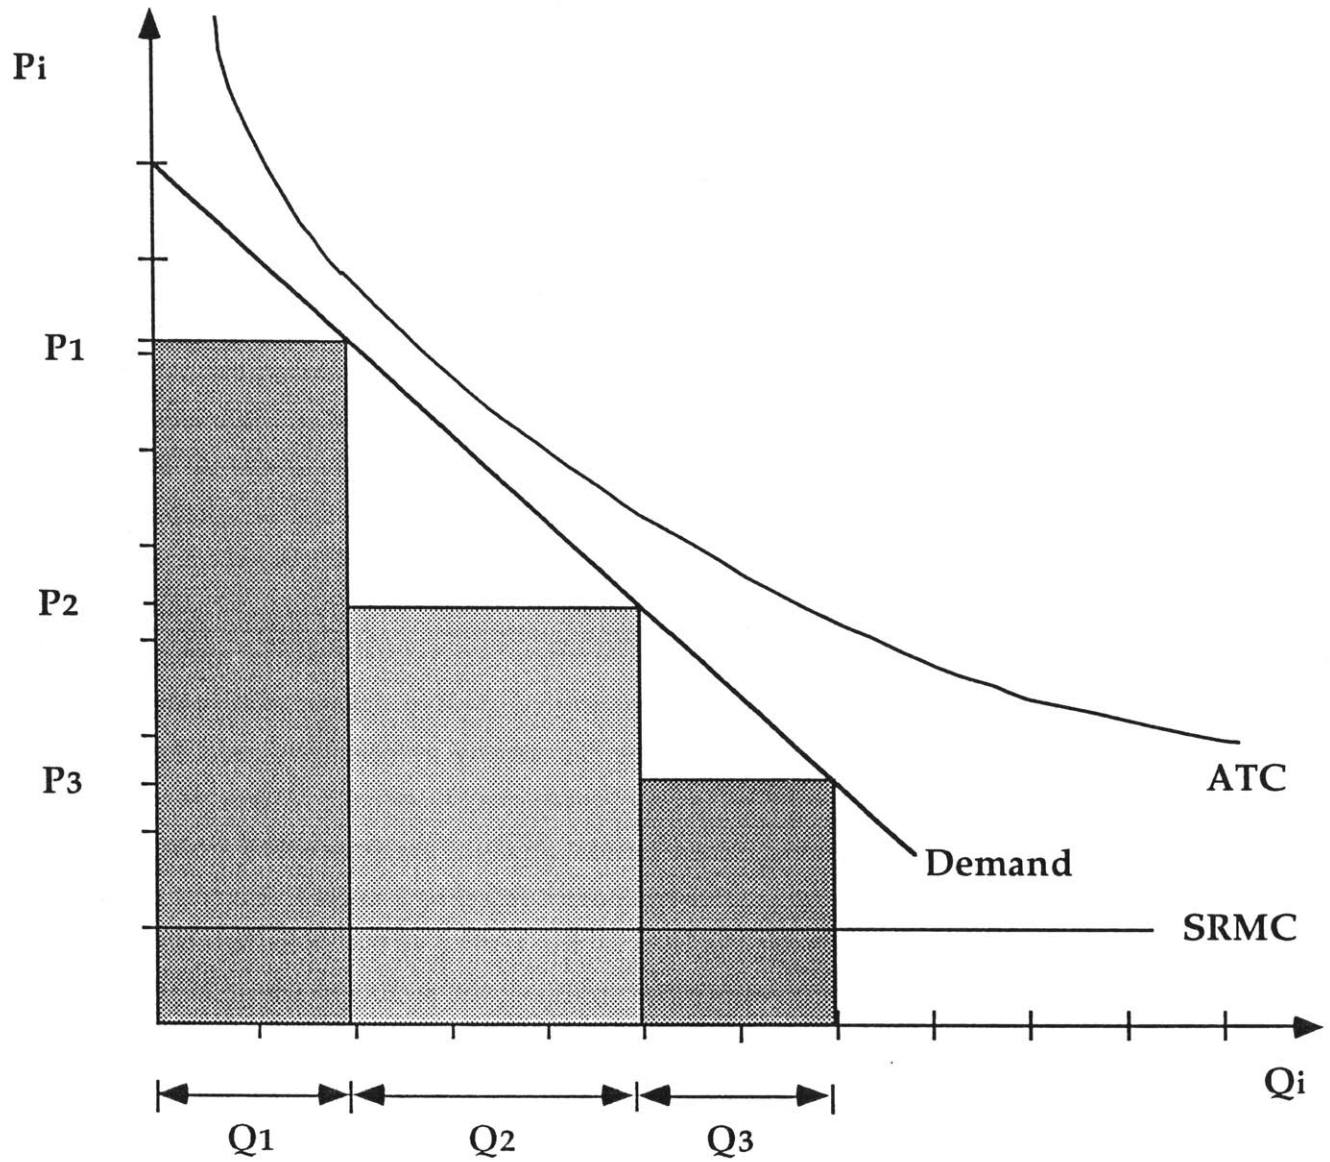
Figure 2.1 - Monopoly Price Discrimination - Three Fare Product Example

and segment the consumer population by their willingness to pay. Thus, the segmentation has been assumed to occur with no cost to the consumer of accepting the additional purchase restrictions. In other words, a seat on the aircraft is viewed as identical no matter what level of inconvenience was endured by the passenger to get that seat at that price level. This assumption also goes against the structure of airline fare products currently offered in the market since advance purchase requirements and Saturday night stay requirements are bound to inconvenience certain low willingness to pay passengers. In reality, the purchase restrictions that are attached to the fare products in the lower classes do not, in fact, perfectly identify and segment passengers by their willingness to pay with no cost. A non-negative cost is imposed on each passenger that is forced to accept an additional purchase restriction.

One of the underlying assumptions of the above price discrimination model is that passengers are all demanding seats on the same aircraft and view that products that are offered to them as commodities. In other words, the passengers do not care about the restrictions that are attached to their fare products as long as they have a seat aboard the departing aircraft. The fact is that airline fare products are not viewed as perfect substitutes, but rather as separate products with common characteristics. Thus, index sorting is not a correct characterization of the majority of price discrimination practiced by airlines.

The characterization of demand found in the price discrimination model presented here would be appropriate if the airline was able to perfectly identify and segment subgroups of the passenger population by their willingness to pay for air travel. Unfortunately for the airline, this ability is unachievable in addition to being in violation of U.S. law (Robinson-Patman Act of 1964). To approximate the desired segmentation of demand, airlines have imposed purchase restrictions and other devices to act as "fences" to prevent passengers willing to pay large amounts for air travel from purchasing discount tickets.

While the assumption about a monopoly market is not true in general, the nature of pricing decisions makes the results of the model presented here instructive to airline managers. Different price levels for the different fare products will be set only at one level in most markets. The matching behavior of carriers that is

enabled through their contact via the computer reservations systems and the Airline Tariff Publishing Corporation makes the result of airline fare products and price levels appear like a monopoly decision at most times. The output of the monopoly market model presented here would not provide an exact representation of the firms and their interactions but rather an indication of the direction in which airlines should move price levels.

2.6 Review of Product Differentiation in the Yield Management Literature

The view of airline fare products, price levels, and passenger demand in the airline yield management literature has been focused on the belief that the separate products appeal to different demand populations and, as such, do not interrelate in a systematic way. Currently, all operational models of airline yield management rely upon the assumption that airlines have the ability to segment their demand populations into homogeneous and independent subgroups. For a taxonomy of the airline yield management and related literature consult Weatherford (1991).

Airlines would like to believe that the self-selection criteria that they employ result in the segmentation of passenger demand into homogeneous and independent subgroups for the differentiated fare products. That is, the effect on the demand for the full fare product of a change in the discount product price level may be viewed as no different from a change in the price of a loaf of bread. The fare product inferiority hierarchy described in Chapter One demonstrates that only under very specific conditions does this independence occur. The view of airline fare products as independent is clearly untrue based upon what is known about air transportation passenger demand.

The fare war that occurred in the Summer of 1992 demonstrates that the demand for fare products is far from independent. Fourteen day excursion fare products were available for sale on the market during the fare war. The price level for the less restricted seven day advanced purchase excursion fare products dropped below that of the fourteen day advanced purchase product. Needless to say, the fourteen day advanced purchase excursion product was not purchased while there was a rush on the seven day excursion fare products. The inter-

relationships between airline fare products and consumer demand for them must be explicitly considered in any framework attempting to model airline fare product differentiation.

In the literature, some of the models used are (third degree) price discrimination models while others are product differentiation models. The price discrimination models assume no cost to passengers associated with their inconvenience. The product differentiation models assume that the demand populations are independent and the demand for each of the fare products is independent of the price levels associated with all other fare products. The appropriateness of the assumptions remains to be addressed in the following chapters of the dissertation.

2.7 Chapter Summary

The chapter has highlighted the modeling of airline fare product differentiation in the context of the economics price discrimination literature. It is apparent that the models that have been used to describe airline markets either have faulty assumptions or do not correctly model the situation. This motivates the need to develop a model of airline fare product differentiation that considers the true problem. In the next chapter, a model is motivated that addresses the shortcomings of the literature and provides useful information to airline management.

Chapter Three

Incorporating Price Levels and Restrictions into the Current Framework of Airline Yield Management

Airline yield management models attempt to maximize expected revenues given fixed price levels and purchase restrictions. The demand level inputs to the models correct for certain systematic trends such as seasonality in the data but do not explicitly consider the effect of fare product attributes on demand levels. This chapter addresses the important factors influencing demand and their treatment within the current framework of airline yield management. The inclusion of fare product attributes in the demand characterization would enable the airline to measure the effects of price levels and purchase restrictions on revenue and passenger demand with the goal of improving revenue enhancement. Provided with information about passenger sensitivity to price levels and product restrictions, an airline could more effectively increase profitability by modifying its fare product offerings.

The chapter begins by addressing the major effects on the demand for airline fare products. Next, the theoretical airline yield management models of demand are examined in the context of their application to a dynamic seat allocation. The effects of data constraints on the application of the theoretical models are then examined. The techniques for correcting demand inputs are reviewed and their shortcomings highlighted. More correct dynamically applicable models of passenger demand are then formulated and the computational and data requirements for their calculation are examined. Finally, the requirements for implementation of the improved demand characterizations are shown to be overwhelming under the current state of practice. An alternative solution is motivated for treatment in the following chapter.

3.1 Major Factors Influencing Air Travel Demand

Many factors affect the level of demand for air travel. Airlines can control some of the factors while being unable to influence others. Even among the factors

seemingly controllable by the airline, the nature of airline competition may make direct influence by a single carrier difficult. In this section, the major effects on air travel demand as well as the ability of the airline to influence them are addressed.

3.1.1 Factors Affecting Demand for Air Travel Overall

Many of the factors influencing the overall level of passenger demand for air travel are beyond the control of airlines. OD market demand levels often exhibit structural shifts as a result of national or international events, particularly when consumers harbor fears about airline safety. For instance, safety concerns lowered consumer demand during the recent war in the Persian Gulf. Traffic levels plunged to record lows in the quarters surrounding the war, most likely as a result of passenger fears of terrorist attacks. Terrorism is often directed at airlines because of the high visibility coverage that airline disasters generally receive in the world media. The media attention fuels passenger safety concerns thus amplifying the downward influence on traffic levels.

Maintenance related problems can cause airplane safety incidents even with the strict standards imposed by safety regulatory agencies like the Federal Aviation Administration. Adverse weather conditions and air traffic control mishaps can cause airplane accidents as well. Regardless of their cause, strings of airplane crashes can produce a downward effect on traffic levels. The effect of airplane incidents takes on importance in the measurement of air travel demand when safety concerns influence the travel decisions of the (discretionary) travel population. As with terrorism, media coverage only heightens the concerns of the population.

Increased modal competition can influence air travel demand levels quite significantly. For instance, if high-speed rail were introduced into a short haul air travel market, demand levels would be expected to drop if rail service levels became competitive with air. The extent of competition considered by the airline must be expanded to account for multiple modes, provided that an acceptable non-airline alternative exists and takes OD market share. An airline must be concerned with the degree to which the new modal competition captures intercity travel market share compared to simply stimulating new demand.

Fortunately for the airlines, the effects of modal competition are confined to shorter haul markets. In the U.S., for instance, the short haul markets experiencing intermodal competition are quite limited.

Traffic growth (or contraction) resulting from changing habits or demographic trends may cause changes in the level of demand for air travel in an OD market. Areas of the country with high population growth, for example, may experience increases in air travel resulting from the larger population. OD market traffic growth may also occur from a shift away from manufacturing industries to professional services at the origin or destination city. Although airlines may respond to the increased demand with higher levels of service in the area, a growth trend itself cannot be influenced by airlines. Demographic changes are a relatively long term phenomenon.

The economic conditions prevailing at the origin or destination city can affect the demand for air travel in an OD market in the short term. Clearly, an area that is conducting less business in general is likely to have less need for air travel. Thus, during periods of local recession demand for air travel may drop. The effects on demand of the economy are not limited to local business cycle fluctuations. Poor performance of the economy on the national or international level can also result in reduced levels of air travel. Cyclic variations in the economy whether local, national, or international, can have a major impact on the demand for air travel.

Finally, the demand for all fare products is influenced by the season of the year. Different OD markets display different peaking behaviors depending upon the nature of the travel occurring between them. For example, the traffic between the Northeast and Florida is high during the winter months, particularly during the Christmas and New Years holidays. Traffic drops significantly in Northeast to Florida markets during the summer months as the weather improves in the North. The seasonal demand differences result from the attractiveness of Florida as a leisure travel destination during the winter months and its decreased attractiveness as a leisure spot during the hot and humid summer months. As the above example suggests, the seasonality of the market often is a major factor in the prediction of travel demand, especially in leisure travel markets.

3.1.2 Factors Affecting the Market Share Between Carriers

Other factors influence the overall level of air traffic for the individual carriers. The influences addressed here affect all fare products but have little to do with the individual fare product attributes. Preferences concerning frequency share, spacing of departures, and carrier attributes influence passenger travel choices and thus demand levels. In any OD market, the itineraries that are available to a consumer on an airline also help to determine the level of demand for that carrier. The discussion begins, however, with perhaps the greatest determinants of carrier market share, non-stop service.

Carriers offering non-stop service in a market offer a superior product to passengers purchasing all types of fare products. Non-stop itineraries are subject to shorter travel times and provide a lower potential for delay than connections. The superior level of service offered by non-stop flight itineraries benefits passengers by increasing the value of service. Moreover, since overall level of service can influence travel/no travel choices, an abundance of non-stop flights may result in higher OD market demand. At a minimum, increased demand levels result for carriers providing non-stop service based primarily on travel time improvements.

Demand levels for air travel are also influenced by passenger perceptions about the different airlines. The carrier perceptions revolve around such factors as reliability, safety, level of in-flight service, and frequent traveler bonuses. For each individual airline, the impact of passenger perceptions can weigh heavily on the overall demand for service. For instance, a carrier that is perceived as unreliable or in danger of the service interruptions caused by a strike or impending bankruptcy may experience reduced demand levels on all flights. Passengers have been known to avoid carriers perceived to have reliability problems. In addition, some passengers simply prefer flying on their "favorite" airline and often go to lengths to ensure passage on that carrier especially since the introduction of frequent flyer programs. Carrier perceptions can result in increased or decreased levels of demand for all fare products depending upon their favorability.

91

Carriers having a high frequency share often carry a disproportionate number of passengers since the number of inbound-outbound itinerary combinations available to a consumer depends on the square of frequency share (Simpson, 1982). Passenger gain resulting from high frequency share is expected to happen in two main ways - increasing passenger itinerary flexibility and providing more attractive departure times. First, passengers traveling on fully refundable fare products are offered more departure choices from carriers possessing greater frequency share. If a traveler arrives at the airport well in advance of his expected flight departure as the result of finishing a business meeting early, for instance, traveling on a carrier with frequent departures at his origin may enable that passenger to travel on an earlier flight than expected. Carriers offering only a single daily flight itinerary in an OD market can offer no such flexibility to passengers desiring itinerary changes. Thus, increased itinerary flexibility is expected to increase demand for a carrier.

Passengers also have preferred departure times when considering air travel. The matching of the flight schedule with the preferred departure times of travelers is an important factor influencing overall demand for a carrier. A carrier offering a large frequency share clearly has the greatest possibility of offering travel itineraries that most closely meet the needs of the travel population. In other words, high frequency share (with proper departure spacing) increases the probability that a given carrier offers a flight departure that coincides most closely with the preferred departure time of the greatest number of passengers. Moreover, actually matching capacity with preferred passenger departure times enables the airline to serve the maximum number of consumers desiring air travel with service coinciding most closely with their preferred departure times and may increase demand levels further. Both individual carrier and OD market demand levels should increase with increased frequency.

The fixed schedule under which an airline must operate in the short term coupled with the physical constraints of the aircraft cycle prevent the airline from matching preferred departure times with passenger preferences in all markets at all times. Carriers can influence their frequency share in any market simply by increasing the number of flights that they offer. Cost constraints, however, limit the extent to which the frequency of service in any market can be augmented.

The advent of frequent traveler programs has increased the effect on the demand levels of carrier perception, favoritism, and departure share at the origin. There is now increased incentive for passengers to choose a single carrier with which to accumulate frequent flyer mileage points to obtain rewards. The incentive to focus on a single carrier results from the non-linear reward structure that characterizes frequent traveler programs. A non-linear reward structure provides increasingly attractive awards for greater mileage point accumulation. Increased market power is afforded to carriers who possess a larger percentage of enplanements at any origin resulting from perceived frequent flyer program captivity. In fact, the preferred carrier of many passengers often becomes the carrier for which that passenger is closest to earning a frequent flyer travel bonus. In the extreme, passengers near a bonus award may even travel more to obtain free travel more quickly and stimulate demand further.

3.1.3 Factors Affecting the Demand for Individual Fare Products

Under deregulation, U.S. airlines have been given complete freedom for the pricing and supply of domestic air transportation services. The freedom entails the setting of price levels as well as the application of restrictions. The demand for air travel is influenced by these fare product attributes. Significant changes to the fare product attributes often influence demand extensively.

Airlines attempt to manage the demand for air travel using price level. It is likely that price levels have the greatest influence on passenger demand of any fare product attribute since the extreme sensitivity of many passengers to price drives many air travel decisions. Airlines lower excursion fares to stimulate leisure demand during periods of historically (or systematically) low travel demand when significant adjustments to the flight schedule are infeasible. For instance, airlines offer discounts to discretionary travelers to stimulate demand during low travel periods at yearly intervals. Special deals on travel to Florida and other warm winter weather locations appear each January and February to stimulate demand after the holiday season.

Since airlines have historically acquired aircraft to meet travel needs during peak seasons, excess capacity has resulted when demand is low. Marketing promotions which offer reduced price travel have been used to fill the excess

capacity during the low seasons. Airlines often reduce fare levels to stimulate demand during high travel periods as well. Thirty to forty percent fare reductions on the lowest priced excursion tickets are often offered in late spring to stimulate demand during the summer. Although greater demand stimulation is needed during the low seasons, less popular flight departures during the summer months can be filled with discretionary travelers who help contribute to fixed costs.

Purchase restrictions also have a large influence on the demand for the individual fare products. Similar to price levels, the individual carriers exercise complete control over the level of restrictions that accompany the fare products that they sell. Changes to existing restriction levels can be made within a day in all domestic markets through the ATPCO clearinghouse. Passenger demand can be managed through a tightening or relaxation of fare product restriction levels. Unfortunately for the airlines, changes to the fare product restrictions are generally viewed as complicated and undesirable by the traveling public. Popular confusion about restrictions limits the freedom of carriers to frequently adjust purchase restrictions. Nonetheless, purchase restrictions and their effects can influence traffic levels for the individual fare products.

Supply effects can influence the amount of traffic carried by the airlines. The availability of any fare product can be limited either by the physical capacity of the aircraft or service class cabin. In addition, booking limits applied to the different fare classes or products can prevent passengers from purchasing a fare product which shows no availability in the computer reservations system for a given flight departure.

The availability of seats for a given fare product has no direct effect on the level of demand for air travel. Passengers desiring air travel who know the fare product structure have preferences that are unaffected by availability. The sole way in which availability influences the actual demand for air travel on any of the individual fare products is passenger discouragement resulting from assumptions about availability limits. For example, a passenger who has previously been denied a booking for his preferred fare product may be discouraged from purchasing air travel. This effect, however, is likely to be very minor. Whether or not the individual fare products are available for purchase by

the consumer is another story. The realization of demand, or the actual number of passengers flown, is strongly influenced by availability. Availability or supply constraints cause realized and actual demand to diverge.

The effect of availability on realized demand becomes important when realized demand data are used as a proxy for actual demand data. The flight loads and passenger counts do reflect the effects of availability on the realization of demand. Binding availability constraints lower the realized demand for all fare products because of the fixed capacity that is provided by any aircraft. When a flight reaches its maximum authorized booking level, for example, all future passenger requests are denied. On the highest demand flights, it is likely that at least some passengers desiring service have been denied because of the capacity constraint and the timing of their reservation requests. The rejected bookings can occur for any fare product. The only way that a rejected consumer appears in the realized demand is through the purchase of an alternative fare product. Thus, although not directly affecting actual demand, the realized demand levels are strongly influenced by availability.

3.1.4 Limitations of Airline Control Over Price Levels and Restrictions

Although airlines as a whole are ultimately responsible for airline fares as well as purchase restrictions, each individual airline exercises far less control over these matters. While any carrier has the power to unilaterally change the price levels and fare products that it offers, the nature of competition prevents carriers from actually pursuing such renegade strategies. Operating in oligopolistic markets, competitive pressures drive carriers to pursue matching strategies for the majority of fare products sold. The mutual dependence among carriers and fears of price wars that characterize the oligopolistic airline industry drive carriers to pursue matching strategies (Wells, 1984).

Strategies aimed at maintaining market share are practiced by the airlines further reinforce matching strategies as the default reaction to change. Moreover, the information contained within the computer reservations systems of airlines make undertaking clandestine strategies in markets nearly impossible. Actions not approved by the other carriers do not often go unnoticed and invite competitive reaction if not reprisal. In the end, the fare product changes that a carrier may

defend in any market depends upon the market strength of the carrier as well as the nature of the change. It is likely that a carrier instituting a fare change that is not matched by the other carriers will withdraw that action.

The ability of an individual airline to unilaterally influence the price levels is generally limited to price decreases. Virtually all airlines will match a lower fare posted by any carrier (at least a major carrier). Carriers are generally unwilling to sacrifice market share to another carrier by allowing a lower published price level in the market. Thus, a carrier can post a lower fare and expect the other carriers to match for fear of being uncompetitive and losing market share. Only in markets where a single carrier exercises considerable market power are fare differentials likely to remain. Unilateral fare increases, on the other hand, are not always matched by the other carriers. Carriers are generally willing to allow other airlines to post a higher fare without responding competitively because of the increased market share expected to result.

Changing purchase restrictions faces resistance from other carriers as well as from consumers. Consumers generally resent the complexity of the airline fare product structure and, as a result, changes to the format with which passengers are accustomed are not often made. Fare product complexities can cause discontent with airlines and adversely affect demand and, thus, are avoided. Most carriers are aware of the passenger sensitivity to fare product complexity and discourage fare product changes through the ATPCO fare product clearinghouse. Competitive forces generally prevent fare product restriction changes. As a result, differences in the airline fare product structure are not widespread with the exception of special fare offerings.

In summary, carriers can influence the pricing and restrictions prevailing in the market. The influence a single carrier has upon fare product prices and restrictions is limited to suggestions to the other carriers sent through ATPCO. Carrier fare product suggestions can be input in the form of either price level or restriction changes. The absolute control that any one carrier can exercise over a market depends upon the level of competition and the willingness of other carriers offering service in that market to match the decisions suggested by the influencing carrier. The importance of remaining competitive drives final airline fare product attribute decisions. The extent to which a single carrier can

influence the prevailing conditions in the market is not outwardly clear and varies by market.

3.1.5 Passengers Book at Different Times Prior to Departure

In addition to information about the factors affecting demand levels, airlines also have interest in the factors influencing passenger booking behavior. This subsection identifies the major factors influencing the incidence of passenger service requests. Flight schedules and associated fare products are posted in the computer reservations systems of airlines over 300 days prior to departure in most cases. As a result, the booking process can begin far in advance of departure and prior to the majority of passenger inquiries. Passengers inquire about travel at different times throughout the booking process. The arrival times of passenger inquiries depend upon several factors including purchase restrictions, ability to plan travel in advance, perceived and actual seat availability to name just a few. The three main areas that affect the passenger booking choice are fare product attributes, trip specific considerations, and passenger attributes.

The attributes of a given fare product can directly influence passenger booking behavior. Fare product restrictions may force a passenger to arrange travel before a certain date. For example, passengers purchasing 14 day advance purchase, non-refundable excursion fare products must book at least 14 days prior to flight departure. Conversely, passengers who purchase non-excursion fare products as a result of travel itinerary constraints may delay booking until just prior to departure if they know that no advance purchase is required. Certainly, the booking behaviors of passengers purchasing unrestricted fare products differs from those purchasing excursion fare products.

Another fare product attribute that affects booking behavior is seat availability since the computer reservations system inventory control structure ensures that each increasingly restricted fare product has less availability than every less restricted fare product. Availability influences the timing of booking decisions based upon the situational considerations of the trip as well as the experiences and characteristics of the individual passenger. For instance, concerns about

97

future fare product availability may induce passengers to secure their reservations early.

Concerns about fare product availability increase when a trip has particular importance to an individual. More-restricted fare products are allocated lower levels of availability resulting from the nested seat allocations practiced by airlines. The reduced availability of lower-priced fare products may encourage passengers to book earlier when planning an important trip. Conversely, a trip that is surrounded by a great deal of uncertainty may not be booked early as a result of the fixed plans not being known until late in the booking process. In this case, the passenger may be more willing to compromise fare product choices in lieu of availability considerations.

Passenger characteristics may influence the booking behavior of an individual. For instance, highly risk averse passengers may book earlier to ensure availability. Risk taking passengers, on the other hand, may delay booking on a particular flight until near departure. The experience of a particular traveler may also influence booking decisions. A passenger who knows that a flight is quite popular as a result of traveling on that flight frequently may book earlier to help ensure a seat on the aircraft. Earlier bookings are expected from passengers who have been denied availability of a preferred fare product recently. The less seasoned traveler, on the other hand, may not know to worry about flight availability considerations and may book later as a result. Passengers not arriving early during the booking process may be denied their preferred fare product or travel itinerary on high demand flights.

In conclusion, the passenger arrival process must be modeled in addition to the other factors that influence passenger demand levels to correctly characterize airline passenger demand. Passenger characteristics, fare product attributes, and trip-specific considerations affect both fare product demand as well as passenger inquiry times. Having highlighted many of the major influences on demand levels and passenger bookings, a look is now taken at the current practice of demand calculation by airlines. The next section focuses on the treatment of demand levels in airline yield management theory and practice.

98

3.2 Current Practice of Airline Yield Management Demand Models

This dissertation seeks to incorporate the price levels of the fare products in the calculation of the passenger demand levels. The expansion of the demand forecasts to include price level would allow the airline to test the impact of price on demand and revenue. Understanding the relationship between price and demand may help the carrier to better increase revenues. An improvement of the inputs to the airline yield management models resulting from better demand forecasts may improve the results of the revenue enhancement techniques. In fact, it has been estimated that a 10% improvement in the yield management system demand forecasting ability will improve expected revenues between 0.5% and 3.0% on high demand flights (Lee, 1990). Clearly, increasing revenue potential holds great interest to all carriers.

3.2.1 Yield Management Model Theoretical Assumptions

In the standard airline yield management calculation methodology, three main inputs are made to the model - the overbooking calculation, the revenue level calculation, and the demand forecast. The yield management system returns the suggested fare class (product) booking limits that are thought to best maximize revenues. The focus of this section is on the demand forecasting portion of the yield management system inputs. The section begins with a brief description of what is theoretically calculated by the yield management model demand inputs. Then, it is shown that many of the assumptions advanced in the theory of yield management models are violated by the practical applications.

Review of Airline YM Demand Calculations

The calculation of booking limits made by airline yield management models requires the distribution of demands by fare class to be provided as an input. The demand forecast is assumed to represent a correct snapshot of the actual distribution of passenger demand expected to face a carrier for a particular flight departure in the future. In theory, the demand levels used as inputs to yield management models are assumed to be known with error. To account for the random error component of demand, instead of providing a single estimate, a probability distribution is calculated for each fare class. The incorporation of

random error into the demand estimates is expected to make the results of the yield management revenue enhancement techniques more robust.

Since the models used in practice are probabilistic, some error not captured by correcting for systematic variations in demand levels is assumed to exist. Unfortunately, the more systematic factors that are omitted from the demand calculation, the more likely that the random error component will be large. In this way, the omission of important systematic factors from the calculation of demand biases the results. Thus, it is important to remove the systematic variations such as those described in the previous section when deriving demand distributions.

A specific distributional form of demand is assumed as input to all airline yield management models. The exact form of the distribution depends upon the preferences of the airline employing the yield management system. Both Gaussian and gamma distributions have been suggested in the literature (Belobaba, 1987 and Smith and Penn, 1988). The choice of distribution depends upon the assumptions made about the true form of demand. With a flexible demand distribution, however, the bottleneck operation remains good quality data inputs for estimation. Favorable revenue improvements have been obtained with both gamma and Gaussian distributed demands.

For a static application of yield management techniques, historical demand levels are assumed to be available from the moment bookings begin until flight departure. The level of demand that occurs for each flight is kept as data for use in the calculation of future demand levels. Using the historical demand data, a distribution of future demand is fitted based upon the distributional assumptions made. The fitted distribution is used as input to the airline yield management model. Demand is characterized by the distribution calculated using the historical data having corrected, of course, for the systematic factors influencing demand level. A more complete discussion of the corrections made to the demand figures is postponed until later in the chapter.

Independence Assumption and Evaluation

Theoretically, the demand forecasts used in airline yield management revenue enhancement models assume the independence of all fare products. In other words, the expected demand level for an individual fare product is uncorrelated with the demand levels for any other fare product. What is known about the hierarchy of fare products leaves the assumption of independence of price level and demand tenuous at best. A more detailed look at the independence assumptions reveals their fallibility in modeling the problem.

The standard independence assumptions surrounding airline fare product demand require that the demand populations for the individual fare products be distinct and separable (Belobaba, 1987). First, distinct and separable populations require consumers to view the fare products as truly differentiated. Since the underlying attributes of each fare product are identical (*e.g.* air travel between an origin and destination at a specific time), the price differential between any two fare products must be representative of the restrictions accompanying them. In other words, the more-restricted fare products must be priced above the less-restricted. In addition, the increased restrictions placed on the lower-priced fare products must be effective (*e.g.* impose a cost on the more price-insensitive consumers). Clearly, no logical consumer would purchase a more costly fare product if the attached purchase restrictions did not adversely affect him.

As with any demand calculation, the attributes of the differentiated fare products must remain stable in the perceptions of the consumers so that no unexpected shifts in passenger preferences occur. The independence of fare product demand distributions is likely if the product remains stable over time. Stability is required, however, in both price level and fences. The relative price levels between the fare products should be stable over time to ensure no shifting of demand over time. This is required since the fare product price levels are not used in the calculation of demand forecasts.

Two of the fundamental components of demand predictions have been omitted from consideration in the current framework - own price and cross price elasticities. In other words, the yield management model demand forecasts assume that the price levels of the fare products have no more effect on

passenger demand than the price of bread, for instance. The effect on the fare product of a price level change in another fare product is limited to the income effect. Given the nature of airline fare levels and their large fluctuations, the zero own price elasticity assumption is rather heroic. In an attempt to correct the demand forecasts, *ad hoc* adjustments are made by yield management analysts before the forecasts are input to the models if large errors in the predicted levels are expected.

Unfortunately for accurate demand forecasting, air travel price levels lack stability and often reason. Carriers do not always set prices at levels that can cover costs under normal circumstances. The price levels that carriers set may be lowered, for instance, to increase short term cash flow by stimulating additional demand. The wide range of fare product price levels that occur based on competitive positioning make forecasting demand independent of price level difficult. The constant posturing by carriers in an attempt to publish favorable price levels in the market results in fluctuations in price levels.

An illustrative example of price level instability is the recurrent industry "fare war" in which a series of discount fare offerings flood the market, usually with little warning. Each fare war significantly drops the price of the lowest cost fare products and results in a lower average revenue per passenger for all carriers. The public response to the drastically lowered fares occurring during an industry fare war is increased demand, demand that is not anticipated by forecasts based entirely upon uncorrected historical data. The surge in demand, however, often greatly exceeds prior expectations and represents a temporary shift in demand. As a result of the wide fluctuations in price level routinely occurring in the air transportation industry, measurements of demands are not always accurate.

In neoclassical microeconomic terms, the stability of fences is required to prevent shifting of the cross elasticities of demand between the fare products. Fence stability ensures that the fare product attributes do not change either outright or in the perceptions of the consumers. In this way, stable purchase restrictions do not necessarily guarantee stable fences if the ability of passengers to meet purchase restrictions changes over time. Thus, the stability of fences is required to the extent that it preserves (reinforces) the independence of the fare product demand distributions.

The demand distribution inputs to the yield management model would need to be recalculated if the restrictions or their effectiveness to segment demand were to change. Fortunately for the airlines, the purchase restrictions accompanying the fare products offered on the market are relatively stable over time. Passenger aversion to complicated fare structures and changes in the service offerings help to force airlines to maintain a relatively stable fare product mix. It is not as clear whether passengers are more able to meet the purchase restrictions over time. However, since the basic fare product structure has remained relatively stable over the past several years, it would appear that airlines perceive no significant shifts in passenger ability to meet restrictions.

The inferiority hierarchy tends to cast doubt on the zero cross price elasticity assumption because of the inherent similarity of the fare products. It is quite likely there is a more significant effect on the demand for E14NR fare products for a change in the price of an E7NR fare product on the same flight itinerary than for a change in the price of bread especially in the presence of booking limits which may close off the E14NR fare product to sales prior to the E7NR, bringing the assumption into question. In the absence of effective fences, for instance, the fare products are identical in all ways except price level. The inappropriateness of the zero cross elasticity assumption can be mitigated, however, with stable, large price differentials that mirror fence effectiveness.

The validity of the independence assumption need not rely upon the belief that passengers have no utility for any product other than the one that they purchase. On the contrary, passengers may prefer different fare products depending upon price levels. The prevailing price levels and interrelationships of the price levels must be stable and constant to justify independence. Unfortunately, there are often wide variations in price level due to competitive pressures and recurrent industry fare wars making the assumption of stable price levels highly questionable. Clearly, assuming the independence of demand levels is not realistic. Under certain conditions, however, the fare product independence assumption has validity. For this to occur, however, the fences between fare products must be effective, the price levels must be significantly differentiated, and passengers must exhibit stable purchasing behavior.

It is interesting to note that the demands used in airline yield management models do not assume knowledge about the willingness to pay of passengers. In other words, no demand function is incorporated in the models. The assumptions about demand levels are made entirely independent of price level. Since the majority of passengers are sensitive to price level, the models implicitly assume an existing and constant price level for the fare product since a stable distribution of demand is assumed to be available. It would be virtually impossible to calculate the demand for a product without either knowing or making assumptions about the product attributes. Thus, at a minimum, the assumption of price level stability has been made about each of the fare products prior to the calculation of demand.

Modeling the Dynamic Process

A dynamic application of the airline yield management revenue enhancement module can lead to an improved solution for the carrier. The state of the art in the modeling of the passenger arrival process used in airline yield management assumes arrival distributions for each fare product between the current time period and flight departure. For example, a demand distribution is calculated for passengers in each fare class expected to arrive during the period between seven days prior to departure and departure itself. In this case, day seven represents a yield management system checkpoint. The bookings to come between each checkpoint and departure are the basis for the calculation of the demand distributions.

The booking limit calculations are updated at the different checkpoints before departure. The demand inputs take account of both the bookings on-hand (already in the reservations system) as well as the forecasts of future expected demand. To calculate the updated booking limits, the yield management revenue optimization model is reapplied to the problem given the bookings on hand and the demand forecasts of bookings to come in every fare class. The seat allocation optimization is run at each of the selected revision checkpoints and new, updated booking limits are imposed The distribution of bookings to come in the period prior to departure, as in the static case, are assumed to be calculated directly from the historical demand data. For a more detailed description of the

104

dynamic application of a seat inventory control booking limit model consult Belobaba, 1987.

3.2.2 Yield Management Input Data Available : Practical Considerations

The data constraints facing airlines implementing yield management systems force some concessions to the theory. In this subsection, the ramifications of the departures from yield management theory are explored. Attention is focused on the passenger demand and booking level data currently used in most yield management models. While the methods are quite effective in practice, a clearer understanding of the assumptions of yield management models and what is actually provided by the demand data may lead to an improved demand characterization. This, in turn, may result in yield management model performance improvements.

The demand forecasting modules currently employed by most airline yield management systems use historical estimates of passenger traffic levels to predict future departure loads. The demand inputs to the airline yield management models are taken from the data stored in the computer reservations systems. The CRS keeps information about the number of passengers who have booked in a fare class and when the bookings occurred. Both sets of information are used as inputs to the dynamic application of the airline yield management models. The demand levels are calculated for the entire booking period of the flight and broken down into the demand arriving between the prespecified checkpoints at the level of detail required for modeling the dynamic process. Unfortunately, there is a discrepancy between the number of bookings on the computer reservations system and the actual flight demand.

Passenger requests for service are often rejected as a result of capacity constraints and booking limits. As a result, the information stored in the computer reservation systems by the carriers only represents realized demand levels not actual demand levels. Thus, the data required by the yield management systems is not actually provided to it. Rather, a proxy measure, the realized demand, is statistically unconstrained and used as the actual input to the yield management model. The extent of the problems associated with using realized demand data in place of actual demand data are addressed next.

105

The airlines are currently limited to the booking and ticket count data which contains the realization of demand subject to supply effects (availability at both the fare product and aircraft capacity level). The demand data collected from ticket counts taken by the airlines and the Department of Transportation represent an estimate of demand constrained by the fixed capacity levels occurring in the OD markets. Only those passengers who found a fare product available appear in the realized demand. There is no guarantee, however, that those passengers received passage on their preferred itinerary or with their preferred fare product. The historical ticket count data simply provides a head count of the passengers and the price that they paid for air travel.

Oftentimes, passengers purchase fare products that are not their preferred ones as a result of availability constraints. When a passenger is denied a request for a fare product on a given flight itinerary, the possibility exists for that passenger to be reaccommodated. The passenger may find a different fare product on his preferred flight itinerary or, conversely, his preferred fare product on a different flight itinerary. For instance, passengers finding no availability for their preferred fare product may purchase a less-restricted fare product, or "sell-up" to a higher fare class. The incidence of passengers being forced to purchase a higher-priced, less-restricted fare product as the result of booking limits placed on their preferred fare product demonstrates the interdependence of the realized demand.

The interaction of supply and demand is ignored when the realized demand is assumed to represent actual demand. Limited supply and the application of booking limits may artificially force passengers to consider the purchase of fare products that are not their preferred ones. The realized demand levels do not reflect the choice process engaged in by passengers, but rather, only the incidence of their final travel decisions. Thus, although a passenger may have wished to have been a member of the demand population for one fare product, availability may have directed him to purchase another.

Passengers who have been denied their preferred fare products do not appear in the realized demand, except perhaps if they decided to purchase another type of air travel service instead. This type of passenger reincorporation violates the fare product independence assumptions that have been made in the theoretical

models. Without binding capacity or booking limits, the realized demand should closely approximate the true demand. However, although there does exist significant overcapacity in the airline industry, the strong peaking behavior of flight departures makes the assumption of no capacity constraint quite dubious. In fact, the number of passengers purchasing a product is always less than (or equal to) the actual demand level for a flight. Thus, realized demand data systematically underestimates the level of demand for air travel with a binding capacity constraint or lower class booking limits since only those passengers finding availability appear.

Passenger bookings on flight itineraries are often canceled prior to departure and prior to purchase. Thus, the booking data that exists in the system is not necessarily an accurate representation of even realized demand but rather realized bookings at a point in time. Because of the refundability of purchased tickets and the non-binding nature of reservations, airlines use passenger "no show" estimates to more correctly characterize expected realized demand levels. Passenger no shows may be the result of a passenger with a fully refundable ticket who did not cancel his booking prior to departure or passengers making bookings may simply not purchase the ticket associated with the booking that they made. Reservations systems often place purchasing deadlines which attempt to remove some of the bookings that do not result in traveling passengers. Even with a paid ticket, passengers may simply decide not to travel and absorb the monetary loss.

The revenue recovered by any passenger is not a measure of his maximum willingness to pay as would be assumed in a standard demand function but rather the result of how much was actually paid considering discounting, ability to meet purchase restrictions, and supply-demand interactions. Although we know that each passenger was willing to pay the amount of his purchased fare product, it is unclear what the maximum willingness to pay (or reservation price) of that passenger was. Thus, the demand distributions simply represent expected passenger counts rather than indications of passenger characteristics.

The demand forecasts used by most airlines use recent historical data for use as input on a rolling horizon basis. In other words, the demand data used for the current forecast is taken from the most recently available data on the system.

Thus, if a structural change occurs in the market the data accessed by the system does not reflect the change. Over time, however, the historical data change to more accurately reflect the current environment provided that a new level of structural stability is reached.

The forecasting of passenger demand is commonly done by flight leg at the fare class level. One underlying assumption affecting demand forecasts is the matching of the fare level with the fare class. For instance, if an E14NR fare product were sold in Q class over the historical demand data period, the accuracy of the future demand forecasts for Q class should depend upon whether E14NR fare products continue to be sold in that class. In the absence of a perfect match between a single fare product being sold in each fare class, there must be some stability of fare products within each fare class to accurately measure the demand populations in each fare class.

Any change to the fare product or mix of fare products being sold in a fare class may bias the forecasting estimates. Since passenger have different sensitivities to fare product attributes, a change in the price level or restrictions associated with the fare products sold in Q class may cause demand forecast inaccuracy. The performance of the yield management system is expected to improve with increased stability in the matching of fare products with fare classes.

The true demand can only be known from actually interviewing the consumers requesting service at the point of sale. A more accurate representation of actual passenger demand might be achieved conducting a telephone survey of passengers who called a travel agent, for instance, to reserve airline travel. Such data collection would require conducting an interview of passengers during the purchase process, something than no airline can do on a large scale basis. In fact, airlines cannot easily conduct such surveys independently since the only truly relevant tests could be performed at travel agencies where passengers are able to choose all available air travel alternatives.

The passenger service requests would represent actual demand if the passenger were given a selection of fare products and flight itineraries and required to choose among them. The data requirements for conducting such a test would be large for even a single origin to destination market. The cost of conducting a

survey for individual markets would be prohibitive. Clearly, aggregation techniques could be used if such a survey were conducted on a representative sample of markets to obtain estimates of demand for similar markets. However, privacy issues as well as travel agent aversion to the administration of such surveys to their customers prevent this type of passenger inquiry from being readily performed. As a result, the true demand levels existing for flight levels are never uncovered and the airlines must rely on proxy measures to predict demand levels using historical realized demand data. There are, however, correction techniques used by airlines that attempt to improve the forecasting results obtained from the realized demand data.

Many of the factors affecting passenger demand levels should be corrected to place more confidence in the demand forecasts used in the yield management systems. The problem with the correction of the factors is the non-quantifiability of many factor influences as well as the large data requirements of the more appropriate demand formulations. The different techniques available to correct the major factors influencing demand are explored next. Suggestions for additional improvements of the demand measures follow.

3.2.3 Correction of Factors Affecting Demand to Improve Forecasts

The many exogenous factors that influence travel behavior are assumed not to directly affect demand levels by most (if not all) airline yield management forecasting systems. External data correction techniques are available provided that data on the exogenous factors affecting traffic levels can be quantified. Currently, airline yield management system forecasts are directly calculated from historical data with only a few corrections. The airlines deseasonalize the data as well as unconstraining it before making the demand predictions. These techniques are explained in this section.

To deseasonalize the data, seasonal indices are calculated for each period of analysis based on the historical demand data. The level of detail of the seasonality indices varies by airline based upon data base capabilities, extent of seasonal demand variations, and technical expertise. The demand data used as input to airline yield management models are normalized to an average measure of seasonality so that improved comparisons can be made. The demand for a

given flight departure is then reseasonalized to more accurately represent the expected traffic level.

Passenger demand forecasts also account for the constrained nature of airline demand data. The supply-demand interaction which appears in the ticket count data is partially corrected using data unconstraining techniques. Airlines rely upon the technique of unconstraining estimates of demand to simulate the expected level of true demand based on when the booking class for a particular fare product was closed to further sales.

When a booking class is closed to further sales in the computer reservations system, the point at which the sale closure occurs is maintained by the yield management system data base. A reference table is kept for which the airline predicts the percentage of bookings in the fare class that are expected by each checkpoint. Using the closure date and the reference table, the expected demand level at departure is calculated using an extrapolation technique. Unfortunately, the technique of unconstraining demand levels is merely a stopgap measure. Many of the supply-demand interactions that occur cannot be accounted for by the simple methods currently used since the unconstrained demand estimate still represents an expectation of the realized demand as opposed to the actual demand. Clearly, airlines would prefer to have actual demand data available to them but, as previously detailed, this is not currently possible.

3.2.4 Major Factors Influencing Demand Not Corrected

After the corrections that are currently made to the data have been completed, several of the most important factors influencing demand levels remain uncorrected. Thus, accurate calculation of demand is difficult even after the seasonal variations and supply-demand interactions are controlled. Structural shifts in consumer behavior are reflected in the data and require that careful attention be paid to the period of analysis chosen to ensure that no uncorrected shifts in OD market traffic demand exist. Impacts such as those previously cited should be explicitly considered or corrected to place more confidence in the demand measures. The incorporation (correction) of such measures is often impractical, however, as a result of the non-quantifiability of the measures.

110

The effects of international events or maintenance related failures on passenger demand levels are difficult to quantify and thus difficult to incorporate into demand forecasts on a systematic basis. The effects of the supply shock resulting from an employee strike at a major carrier, for instance, face similar difficulties. The incorporation of this type of exogenous demand influence is quite difficult except on an *ad hoc* basis. As a result, these impacts must be omitted from any systematic measurement of demand until a quantifiable measure is developed that accurately represents them. In the meantime, the revenue management department at an airline can attempt to control such demand shocks on an *ad hoc* basis.

More importantly, the price levels and restrictions of the fare products are not explicitly considered when the passenger demand distributions are calculated. Changes to the fare levels, which are frequent, may greatly influence the demand for the individual fare products, perhaps more than economic conditions or safety concerns. Thus, an operational method for including price level and restrictions in the calculation of demand would increase airlines ability to maximize revenue.

Ignoring the important impact of price level on demand can lead to biased forecasts. Consider (again) the example of the fare war of the summer of 1993. The price level of the E7NR fare product was dropped by 50% to a very low level. The result was a flood of bookings for the E7NR fare product. Clearly the own price elasticity of demand played an important role in the influence on demand levels. Cross price elasticity for the E14NR fare product also was needed to predict the effect on demand levels that occurred during the period. The price of the E14NR fare product was also dropped to a level identical to the E7NR fare product and both services were offered on the market simultaneously.

Recalling the properties of the fare product inferiority hierarchy, the demand for the E14NR fare product should have dropped to zero since the E7NR fare product is superior. Clearly, the absence of the cross price elasticity relationship between the E14NR and E7NR fare products resulted in biased demand forecasts for the two products during the fare war. A correct demand forecast would predict that all of the demand for the E14NR (in addition to all stimulated demand) be absorbed by the E7NR fare product . Admittedly the example

provided by the fare war is extreme, however, it effectively illustrates the need for the incorporation of price level into the calculation of the demand distributions that used as input to the airline yield management models.

3.2.5 Importance of Demand Forecasting Improvements

The independence assumption of the fare products still prevails in airline yield management forecasting. The incorrect characterization of demand biases the demand forecasts systematically. The data constraints facing the airlines only allow calculation of imperfect demand distributions made worse by the wide fluctuations in the levels of price and traffic demand that characterize air travel. Correction of the major factors affecting demand for air travel may improve forecasts significantly.

More complex models that consider a greater number of factors that affect transportation demand can potentially provide an improved method for capturing passenger demand levels than the independent calculations currently performed. Any relaxation of the assumptions made by the current airline yield management demand forecasting modules can lead to improved measurements and, in turn, improved airline profitability. In the next section, the factors influencing demand level are incorporated within the current framework of airline yield management.

3.3 Formulation of Demand Distributions Conditionally Dependent on Price

The calculation of demand distributions can be extended to incorporate the important factors influencing passenger traffic levels by formulating them as conditional distributions. With the proper formulation, the conditional distributions can more correctly represent the expected demand levels for an expected state of the system. It is hoped that the improved level of demand calculation given the current state of the system can be input to the yield management models to produce superior revenue enhancement results. The conditional distributions can demonstrate the effect of price levels on passenger demand. In addition, the ability to input different distributions for different states of the system can allow measurement of the effects of price levels on

112

expected revenue. The formulations conditional on price level appear in this section.

3.3.1 Demand Distributions Conditional on All Price Levels

An ambitious and improved formulation of demand would incorporate all of the price levels prevailing in the market during the current period. Provided that enough observations of demand level were available at the prevailing price level, the distribution conditional upon all fare product price levels would incorporate the own price and relevant cross price elasticities explicitly. The following distributions of demand based upon all fare product price levels could be calculated as :

$$P(D_i) = P(D_i | P_i, P_j \quad \forall j \neq i)$$

where D_i is equal to passenger demand for fare product i and P_i is the price level of fare product i. The calculation of the above probability distributions would require knowledge of the fare levels of all fare products sold in the market during the current period.

It should be noted that although the conditional demand distributions include only price levels and not restrictions, the latter attributes may be incorporated into the models in a similar fashion. However, the stability of the fare product attributes is likely to make the explicit consideration of restrictions in demand calculations unnecessary, especially given the costs associated with maintaining the additional data. Thus, the incorporation of restrictions does not appear here.

The large number of fare products sold on the market and, equivalently, the large number of fare classes used by airline yield management systems would require the measurement of many price levels. Because of the large number of price measures required, the stability of the price levels would become even more of an issue so that enough observations of demand at the current fare levels would exist to calculate the conditional distributions. The odds of securing enough observations for the calculation of such a distribution are very low. As a result, there is a need to simplify the conditional distribution to have a better chance at a feasible estimation.

3.3.2 Demand Distributions Conditional on Own Price Levels

To simplify the estimation, all but the most important price level calculations are left out. The price level of the product itself is the sole factor that is conditioned upon to more correctly formulate demand levels in an attempt to meet the data constraints. Nonetheless, an improvement in the specification would result from the introduction of the important characteristic of own price level into the demand calculation. The influence of a fare product's own price level on demand can be improved over the independence assumption by calculating the conditional distribution of incremental demand with respect to a product's own price level. The distribution conditional on own price level is formulated as follows :

$$P(D_i) = P(D_i \mid P_i)$$

where D_i is equal to passenger demand for fare product i and P_i is the price level of fare product i. The data requirement and length of the forecast horizon would be reduced for the calculation of the simpler model based only on the own price relative to the (all) price level model shown previously. Provided that the data were collected for demand/price level observations occurring in the market, the conditional probability distributions could be calculated from the historical data. Unfortunately, several problems highlighted below still complicate the calculation of the conditional distributions.

3.3.3 Data Requirements

The data for historical fare levels are not stored at present in the yield management data bases used by most airlines. The data requirements for calculation of the probability distributions of demand conditional on price levels are quite large since the fare product structure that characterizes airline OD markets generally has, at a minimum, three selling fares for the coach class cabin alone. Also, the observations of demand levels at the currently prevailing price levels must have been available during the historical time frame maintained by the data base which, in general, is not yet true. Moreover, the price level data must be stored for the different fare classes over a long period.

Maintaining enough observations on each of the available fare products would require a long period of analysis and large amounts of storage space. The tradeoff existing between the number of observations collected and the length of time over which the data are taken must be assessed. Clearly, a large enough number of observations are required to calculate a probability distribution for each state of the system, or prevailing fare product structure and set of price levels.

The period over which data are collected must be long enough to estimate the parameters as well as providing for reasonable forecast accuracy. For a given structural state of the system, the more valid observations the better. The possibility that some type of structural change has occurred in the market increases with the historical time frame used. Simply stated, the longer the period over which the data are taken, the more likely a structural change is to have occurred in the market. In this way, conducting the data collection over a long period of time would increase the likelihood of econometric problems compromising the analysis.

The storage costs associated with retaining historical revenue level data alongside the historical passenger demand data must also be considered. Passenger booking behavior is currently retained over an approximate two month basis. The additional requirements of the price level data acquisition and storage need also be identified and evaluated. As a result of the larger data requirements associated with calculating conditional distributions, the period of historical demand (as well as revenue) data that must be used is likely to increase. The tradeoff between the benefits of improving the formulation of demand and the increased costs (along with the ability to perform a high quality conditional distribution estimation) must be evaluated to assess the benefits of an implementation of conditional demand distributions.

3.3.4 Price Level Measures

To complicate matters further, the measurement of price levels is not outwardly clear. Proxy measures for price levels are available in the same form as the revenue inputs to the yield management models -- average revenue per passenger by fare class. The average revenue per passenger by fare class

115

emanates from a ticket coupon sample taken by the carrier operating the yield management system. Unfortunately, the revenue per passenger figures do not correspond directly with the fare products sold on the market (although the influence of the OD market selling fares on realization cannot be denied). The revenue per passenger figures represent a realization of the price paid instead of the actual price levels facing consumers. Moreover, if the revenue per passenger data are used to calculate the distributions, it may be quite difficult to match the fare levels that were facing the passengers when they purchased air travel with the concurrent fare class booking behavior.

The problems associated with matching booking behavior with the prevailing fare product structure make any effort to estimate demand distributions conditional on price levels difficult. Clearly, the booking behavior that takes place under a certain fare product structure may not be the same when the product mix changes. Large scale changes to the price levels of the fare products even under the same fare product structure may also change the booking behavior of consumers. Thus, the matching of booking behavior and price levels is tantamount to accurately calculating the conditional distributions.

A time lag is introduced in the yield management system resulting from the data acquisition of the revenue per passenger figures. The time lag makes the matching of the price levels prevailing in the market and the current demand levels difficult to obtain for the current period. The average revenue per passenger numbers are not processed until the coupons have been lifted from the passengers and the data have been cleaned and processed.

The matching of revenue per passenger numbers and traffic numbers may be improved with the matching of fare products and fare classes. If a single fare product were associated with a single fare class, the revenue numbers could be obtained from the current selling fares on the market. The best demand forecasts would be constructed if the booking limits were set and the data taken for the individual fare products. The discrepancies between the fare classes would then disappear. However, the "odd-ball" fares that are posted in many markets along with the shifting of fare product offerings would add other difficulties to the problem not encountered with fare class level data.

116

There is difficulty in assigning a single selling fare to a fare class in many OD markets. The calculation of the distributions would require data provided at the most detailed level with price level stability over a significant time period. Stability of price levels would be required to calculate the conditional distributions of the different fare products. This requirement is clearly unrealistic given what is known about the fluctuations occurring for the fare products. Thus, actual implementation of one fare product per fare class would not likely be feasible without a simplification of the system-wide OD market fare product structure. Attempts at simplification of the fare product structure do not appear to be sustainable within the current competitive framework as witnessed by the failed attempt made by American Airlines in 1992.

In the absence of perfect fare class-fare product matching, a weighting of the fare products by traffic level would be necessary which requires an assumption about the distribution of demand within the fare classes. Some sort of weighted average of the fare products within the fare classes might be constructed. The weighted average would allow an immediate calculation of the fare products selling in each fare class. The fares available on the market are available to the airline through the ATPCO data base. Unfortunately, the weighting of the selling fares on the market by within the fare classes that they are sold may be a more difficult problem than estimating the conditional distributions.

The airline would need a distribution of fare product sales by fare class and this measurement would need to be stable over time to provide for an accurate weighting. The lack of stability and large estimation requirements of such a weighting scheme are likely to be prohibitive. Thus, the ability to calculate the distribution of demand by fare product within the fare classes, however, is beyond the current capabilities of the airline data bases.

In summary, the calculation of demand levels conditional on price levels would provide an improvement in the forecasting measures. However, the costs associated with collecting this information coupled with the lack of stability of the pricing data make the calculation of such a model questionable. In any case, the barriers to calculation are too large to enable implementation within the current framework. Thus, an alternative methodology for calculating the effects of price levels and restrictions on air travel demand and revenue is necessary.

117

3.4 Chapter Summary

The chapter began by discussing the factors affecting passenger demand for fare products to provide the motivation for more correctly specifying demand. The assumptions of existing airline fare product demand models were then assessed to motivate the need for a more correct (detailed) calculation of passenger demand. Demand distributions that are conditional upon price level were then formulated. Finally, the data needs and practical requirements for the estimation of the conditional passenger demand distributions were discussed. As a result of the difficulty associated with the estimation of the conditional distributions, the need for a more tractable model for passenger demand was clarified. The static model of passenger demand appearing in the next chapter addresses the fare product dependency relations.

Chapter Four

A Static Generalized Cost Model of Airline Fare Product Differentiation

A model is developed with the intention of gaining insight into questions concerning airline pricing and fare product differentiation. However, since the dynamic process is too complex to model the problem has been simplified by modeling it as a static process. Not only does the static assumption allow for a tractable model, but also provides the ability to test the sensitivity of the underlying relationships between passengers, revenues, price levels, and purchase restrictions.

4.1 Modeling Notation and Scope of the Problem

In order to decrease the complexity of the analysis, the scope of the problem addressed is at the single carrier, single flight level in an isolated OD market. Coach class cabin service is assumed to be offered exclusively so that, although passengers are offered a variety of differentiated airline fare products, the level of in-flight service that each passenger receives is identical.

Each consumer is offered the choice between N differentiated airline fare products for the flight departure. The fare products decrease in value as additional purchase restrictions are added. In the notation, fare product 1 represents the full fare, unrestricted ticket while fare product N represents the most restricted ticket available to consumers. It is assumed that fare product i is more restricted than fare product i-1.

4.2 Modeling the Demand for More-Restricted Fare Products in Isolation

Imagine that the days of U.S. domestic airline regulation had returned and the only fare product available for air travel in the origin to destination (OD) market was the full fare, unrestricted product (referred to as fare product 1). Prior to the introduction of the "Supersaver" airline fare products in the mid-seventies, such

a ticket was the only choice in the coach class cabin for U.S. domestic service. A single product market demand function can be constructed representing the price level/demand relationship for all passengers aboard the aircraft. In this scenario, the number of passengers choosing to travel at any single price level would be :

$$Q = f_1(P_1) \tag{4.1}$$

where

$$
\begin{aligned}
Q &= \text{number of passengers purchasing fare product 1} \\
f_1(\bullet) &= \text{market demand function for fare product 1} \\
P_1 &= \text{price of fare product 1 where } P_1 \in [0, \infty]
\end{aligned}
$$

In other words, the single price level P_1 set for the unrestricted fare product would result in Q units of demand according to the price-demand relationship $f_1(P)$ prevailing in the market.

Imagine next that the regulatory body ordered the removal of all unrestricted fare product sales from the OD market. Instead, all purchases of airline fare products were required to be made at least three days in advance of the flight departure. A market demand function similar to that of the unrestricted fare product would exist for the three day advance purchase (AP3) ticket. The market demand function for the AP3 ticket, called fare product 2, has the form :

$$Q = f_2(P_2) \tag{4.2}$$

with notation analogous to that of the unrestricted fare product demand function.

Although the two fare products offer the same level of in-flight service, the AP3 fare product possesses a less desirable attribute when compared to the unrestricted product since passengers are forced to meet the 3 day advance purchase requirement. This makes the restricted product less valuable to consumers than the unrestricted fare product. In fact, the restricted product can never be more valuable than the unrestricted product. Thus, the market demand function for the more-restricted, less desirable (AP3) product is everywhere

lower than the unrestricted product because of the degradation in product quality resulting from the applied purchase restriction. In fact, for any restricted fare product, consumers experience "degradation costs" associated with accepting additional purchase restrictions. Mathematically, the expected difference in market demand is represented by the relation $f_2(P) \leq f_1(P)$ which holds for all values of price level P.

Isolated demand functions for other fare products can be motivated similarly to the AP3 fare product. A series of fare product market demand functions can be constructed:

$$Q = f_i(P_i) \tag{4.3}$$

each representing a unique characterization of passenger demand subject to increasing (in i) levels of restrictions. The demand for air travel is lower for fare products with more restrictions. Since each fare product i+1 has been defined as more restricted than each fare product i, the demand function for fare product i has the property $f_{i+1}(P) \leq f_i(P)$ for all values of P. Simply stated, there are fewer passengers willing to travel at any price level as the fare product offering becomes more restrictive.

The demand for the more-restricted products, $f_i(P_i)$ for all i>1, over the entire population is somewhat contrived, however, since no rational airline or regulatory body would needlessly constrain or restrict a product without justification. The underlying structure of isolated airline fare product demand is, however, of interest to airlines. Airline fare products are differentiated using purchase restrictions in an attempt to segment the consumer population. The next section introduces a model that combines the passenger segmentation ability of the airline with the underlying demand structure to more correctly characterize airline fare product differentiation.

4.3 Basic Airline Fare Product Differentiation Model Framework

Airlines place restrictions on their fare products in order to identify and segment passengers on the basis of their willingness to pay for air travel. The

121

segmentation relies upon the airline's ability to capitalize on the relative sensitivities of passengers to fare product attributes such as convenience and flexibility. The fences described in Chapter One gave a detailed look into the techniques used by airlines to achieve such a market segmentation.

The model proposed in this section assumes that by designing the purchase restrictions associated with each less-desirable fare product, the airline has the ability to <u>perfectly</u> segment the consumer population into N subgroups given the prevailing price levels P_i for the N differentiated fare products. The subgroups are assumed to be ordered by willingness to pay with their members possessing similar sensitivities to fare product attributes. Although helpful for a more clear exposition of the basic model, the ability of the airline to perfectly identify and segment the demand population is relaxed later in the chapter.

Degradation costs resulting from the passenger segmentation devices are explicitly considered in the modeling framework. The degradation costs that are incurred by consumers only serve to lower the demand in the N-1 restricted fare product subgroups *ceteris paribus*. Including degradation costs incurred by passengers resulting from the imposed purchase restrictions represents a more realistic view of the airline fare product differentiation problem than that provided by the standard monopoly price discrimination model described in Chapter Two (*e.g.* Kahn, 1970). The incorporation of the imperfections associated with self-selection devices in airline attempts to segment the consumer population more correctly represents current industry practice. In this way, explicit consideration of degradation costs extends the airline yield management literature.

The modeling framework of airline fare product differentiation presented here is motivated first using a simple example in which passengers must choose between three fare products. The three fare product example is then extended to model the general N fare product case. In the first example, the three fare products must be designed by the airline to identify and segment the population by their willingness to pay. Fare product 1 (as in all cases) has no attached purchase restrictions. Assume that fare products 2 and 3 have advance purchase requirements of three and seven days, respectively, and fare product 3 also requires a round trip purchase with a Saturday night stay. The assumption of

specific restrictions for fare products 2 and 3 is solely for the sake of expository clarity; fare products 2 and 3 may represent any two restricted fare products. In the three fare product examples presented, fare products 1, 2, and 3 are sometimes referred to as FP1, FP2, and FP3, respectively.

4.3.1 Three Fare Product Model

The model of the three fare product market uses the isolated fare product demand functions to build a representation of OD market demand. The passenger demand populations for the three fare products are determined one after the other according to an assumed passenger fencing structure. The final passenger demand resulting for the OD market is then derived using the fare product differentiation model and a set of fare product price levels.

To correctly specify the model, the isolated market demand functions for the three fare products must have the property :

$$f_3(P) \leq f_2(P) \leq f_1(P) \tag{4.4}$$

for all values of P. In other words, at any single price level, passengers prefer fare product 1 to fare product 2 which, in turn, is preferred to fare product 3. These inequality relations guarantee that the isolated fare product demand functions never cross (although they may meet).

The model begins by assuming that only the unrestricted fare product (1) is offered for sale in the OD market. A price level P_1 is set and the airline expects $Q_1 = f_1(P_1)$ passengers to request a seat, as previously motivated. Referring to the demand function in Figure 4.1, it is clear that the Q_1 passengers possessing the highest willingness to pay for air travel of the entire demand population purchase the unrestricted fare product.

Passengers purchasing the unrestricted product are believed to be extremely sensitive to fare product attributes (*i.e.*, having no advance purchase or Saturday night stay requirement). Historical bookings of unrestricted (high-priced) fare products support this belief. In the model, the Q_1 fare product 1 passengers are assumed to unconditionally prefer the unrestricted product when it is offered at

123

Isolated Demand for the
Unrestricted Fare Product
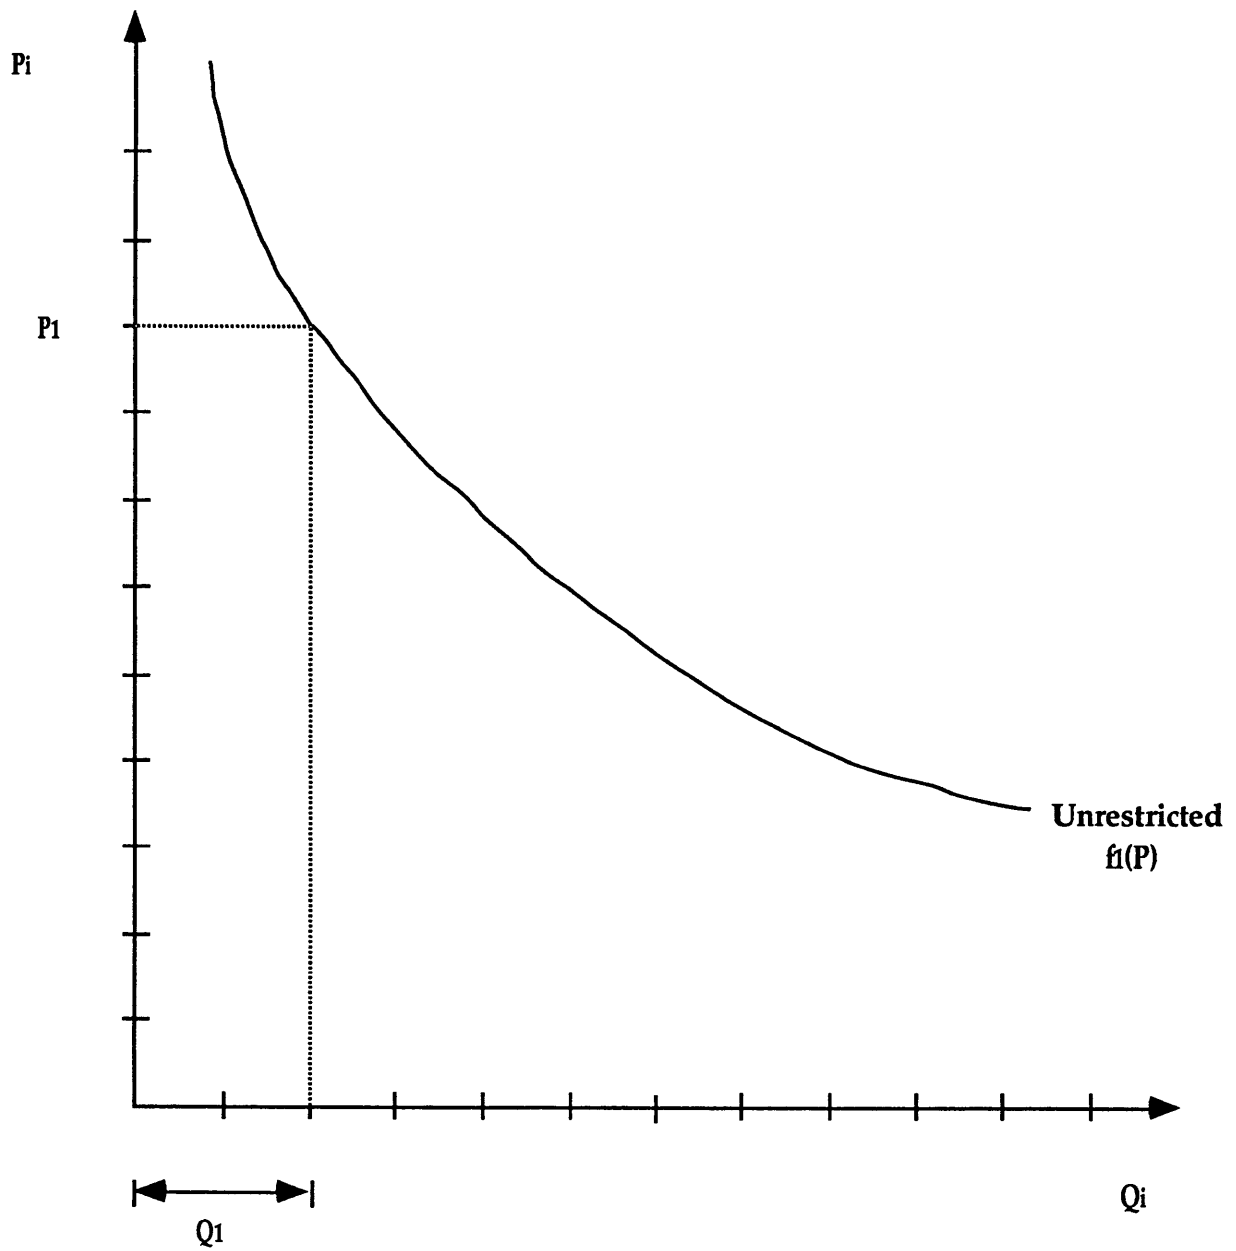
Figure 4.1 - Isolated Demand for the Unrestricted Fare Product

price level P_1. The Q_1 passengers with the greatest willingness to pay are, in effect, assumed captive to the unrestricted fare product. Or, they can be thought of as subject to a perfect fence.

Next, the fare product 2 is made available to passengers wishing to travel in the OD market and offered alongside the unrestricted fare product. The fare product 2 demand function is equal to the portion of the isolated FP2 demand function (similar to the AP3 demand function motivated in the previous section) to the right of the first Q_1 passengers, as shown in Figure 4.2. To prevent a double counting of the individual passengers that make up the fare product 2 market demand function, there must exist a one to one correspondence between the Q_1 passengers captive to the unrestricted fare product and the first Q_1 passengers, in order of willingness to pay, on the isolated fare product 2 demand function. In other words, the passengers willing to pay the most for the unrestricted fare product must also be the ones willing to pay the most for fare product 2.

When a price level P_2 is set, passenger demand for fare product 2 is calculated using the formula $Q_2 = f_2(P_2) - Q_1$. For the airline, an additional Q_2 passenger requests result from the introduction of fare product 2 onto the market. The incremental passengers are willing to travel using fare product 2 at the price level P_2 but were not willing to pay P_1 for the unrestricted fare product. Similar to the Q_1 captive fare product 1 passengers, the Q_2 fare product 2 passengers are assumed to prefer that product to any more-restricted fare product based on their sensitivities to fare product attributes.

Finally, fare product 3 is offered for sale to the remaining passengers in the market. The portion of the demand population interested in purchasing fare product 3 has been reduced by $Q_1 + Q_2$ passengers as a result of offering the fare products 1 and 2 to the more flexibility-sensitive consumers. The demand function for fare product 3 is equal to the portion of the isolated FP3 demand function to the right of the initial $Q_1 + Q_2$ passengers, as shown in Figure 4.3. Setting the price level P_3 will attract $Q_3 = f_3(P_3) - Q_1 - Q_2$ passengers for fare product 3. Again, a one to one correspondence of individual passengers along the three demand functions is assumed.

125

Demand for Fare Product Two
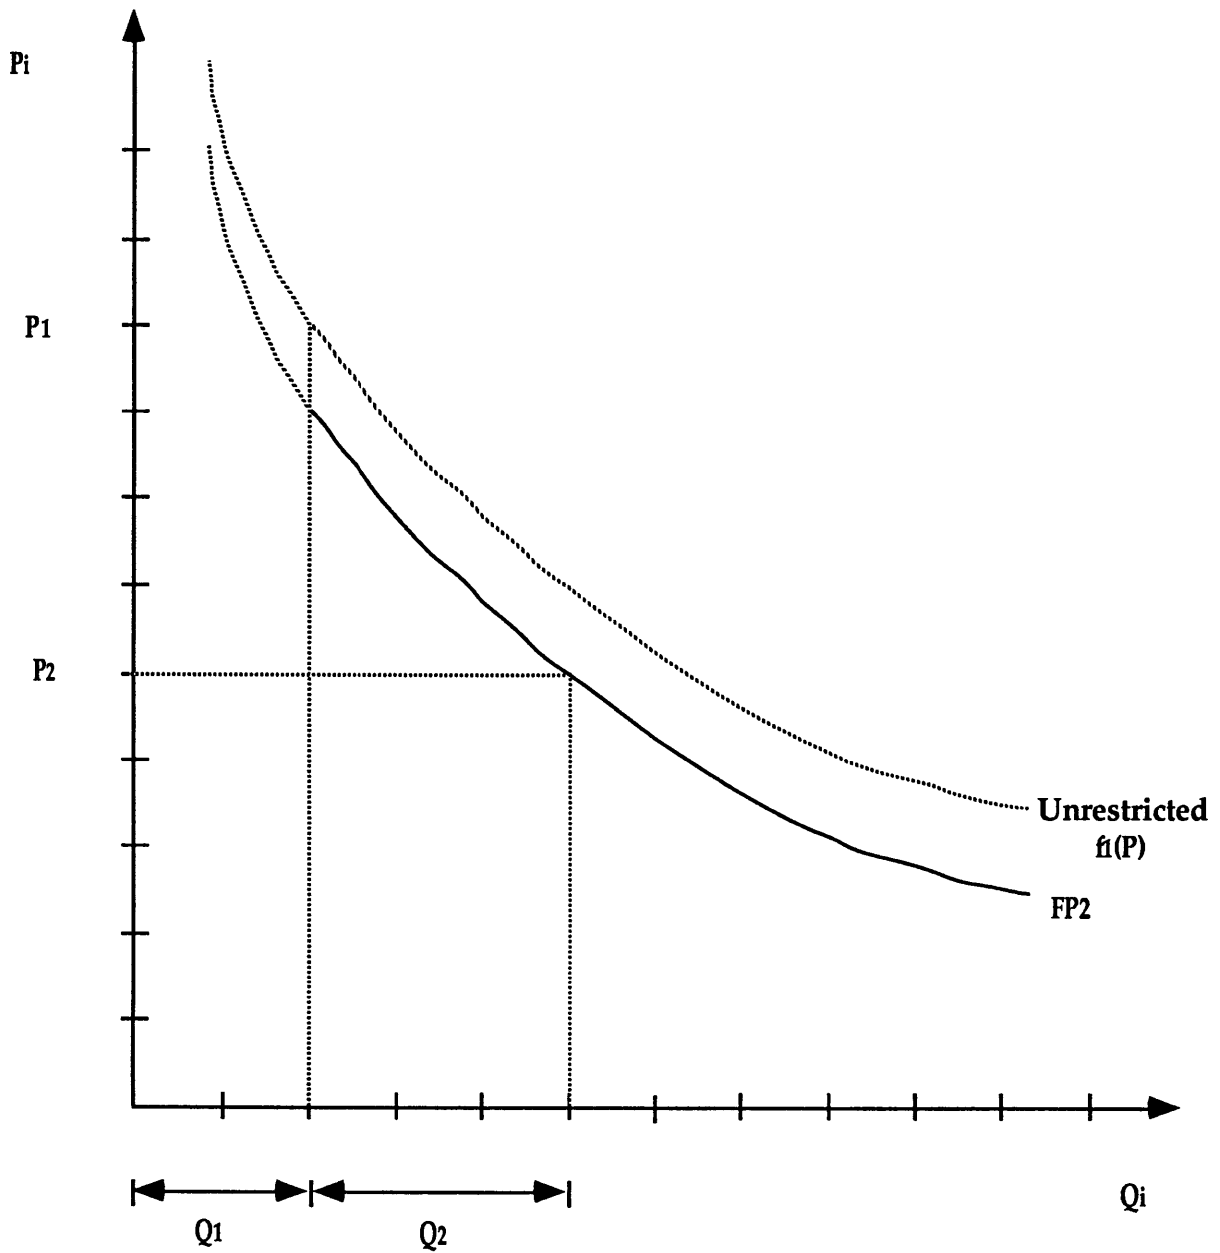
Figure 4.2 - Demand for Fare Product Two

Demand for Fare Product Three
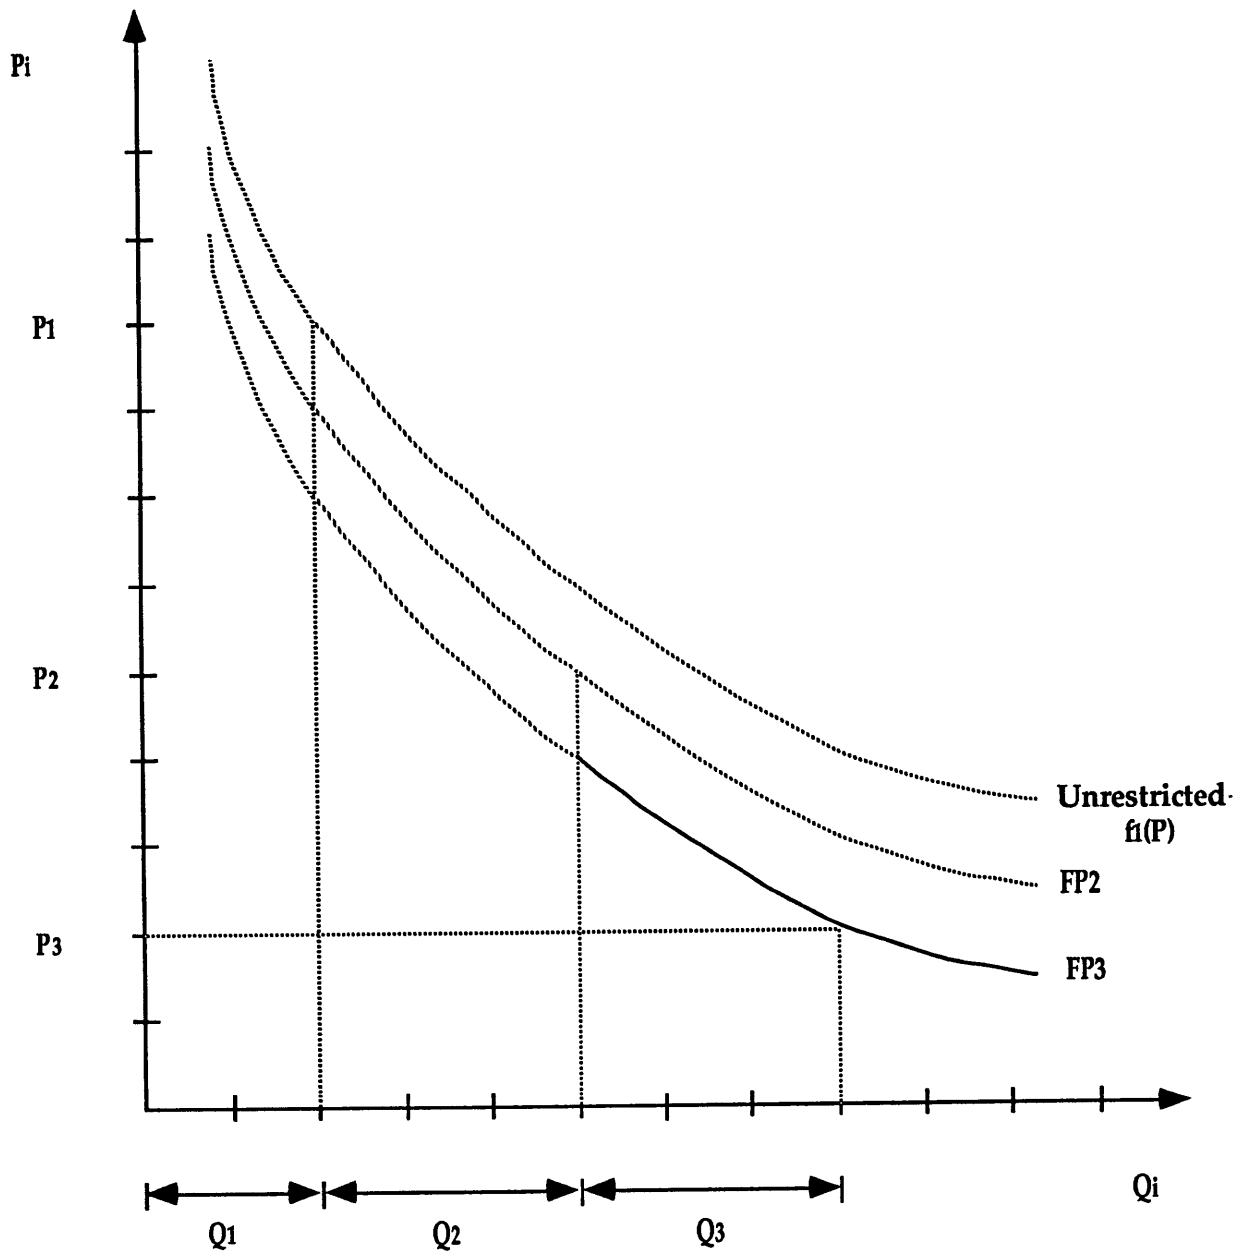
Figure 4.3 - Demand for Fare Product Three

It should be noted that the case in which $f_j(P_j) \leq f_i(P_i)$, where fare product j is more restricted than i, is assumed never to occur. Under that scenario, the demand for the more-restricted fare product j would be less than that of the less-restricted fare product i This would result in no demand for fare product j under the assumed passenger segmentation scheme. Since the passenger segmentation schemes available to airlines are not very fine-grained, however, this case is unlikely to occur in practice. Thus, little generality is lost by ignoring this degenerate case.

In summary, given the price levels P_1, P_2, and P_3, the total demand for the flight departure is simply the sum of the Q_1, Q_2, and Q_3 passengers purchasing fare products 1, 2, and 3, respectively. A representation of passenger demand for the market in the three fare product model appears in Figure 4.4. The revenue contribution from each fare product is represented in the shaded regions of the figure. Moving from left to right in the figure, the 3 shaded boxes represent the contributions of the unrestricted fare product (P_1Q_1), fare product 2 (P_2Q_2), and fare product 3 (P_3Q_3). The $Q_1 + Q_2 + Q_3$ passengers make a total revenue contribution of $P_1Q_1 + P_2Q_2 + P_3Q_3$ to the airline.

Welfare and consumer surplus implications can also be calculated directly from the three fare product model. The consumer surplus associated with each fare product is simply the area between the fare product demand function and the horizontal fare product price level line over the range of passengers purchasing that fare product. Thus, even with a perfect passenger segmentation, some surplus is returned to consumers.

Figure 4.5 shows the consumer surplus associated with each fare product in a three fare product example constructed using isolated <u>linear</u> fare product demand functions (for expository simplicity). The lightly shaded triangle areas labeled CS1, CS2, and CS3 in the figure represent the consumer surplus associated with fare products 1, 2, and 3, respectively. Mathematically, the consumer surplus associated with each of the three fare products is :

128

$$\text{Consumer Surplus}_1 = \int_0^{Q_1} f_1^{-1}(q)dq - P_1Q_1 \qquad (4.5)$$

$$\text{Consumer Surplus}_2 = \int_{Q_1}^{Q_1+Q_2} f_2^{-1}(q)dq - P_2Q_2 \qquad (4.6)$$

$$\text{Consumer Surplus}_3 = \int_{Q_1+Q_2}^{Q_1+Q_2+Q_3} f_3^{-1}(q)dq - P_3Q_3 \qquad (4.7)$$

where the variables are defined as before. The isolated fare product demand functions $Q_i = f_i(P_i)$ are assumed to be invertible.

To demonstrate the inefficiencies caused by the degradation costs resulting from applied purchase restrictions on an OD market, the welfare loss experienced for a given composition of passenger demand by fare product is analyzed. The price levels in the case considering degradation costs are set to levels which maintain the composition of passenger demand for each fare product occurring in the zero degradation cost case. In other words, price levels are modified to carry the same number of passengers (by fare product) when degradation costs are imposed as in the zero degradation cost case. Both revenue and welfare losses result for a constant passenger mix when degradation costs are introduced to the model. The lower value of system welfare from the generalized cost model assuming degradation costs may be subtracted from the welfare value in the zero degradation cost case to provide the deadweight welfare loss (as defined in this dissertation). All subsequent discussions of lost welfare examine the losses incurred by a like passenger mix when degradation costs are incorporated (using a similar methodology).

At the fare product level, the welfare loss associated with each more-restricted fare product is equal to the area between the isolated unrestricted fare product demand function and the more-restricted fare product demand function over the

129

Representation of OD Market Demand
Three Fare Product Example
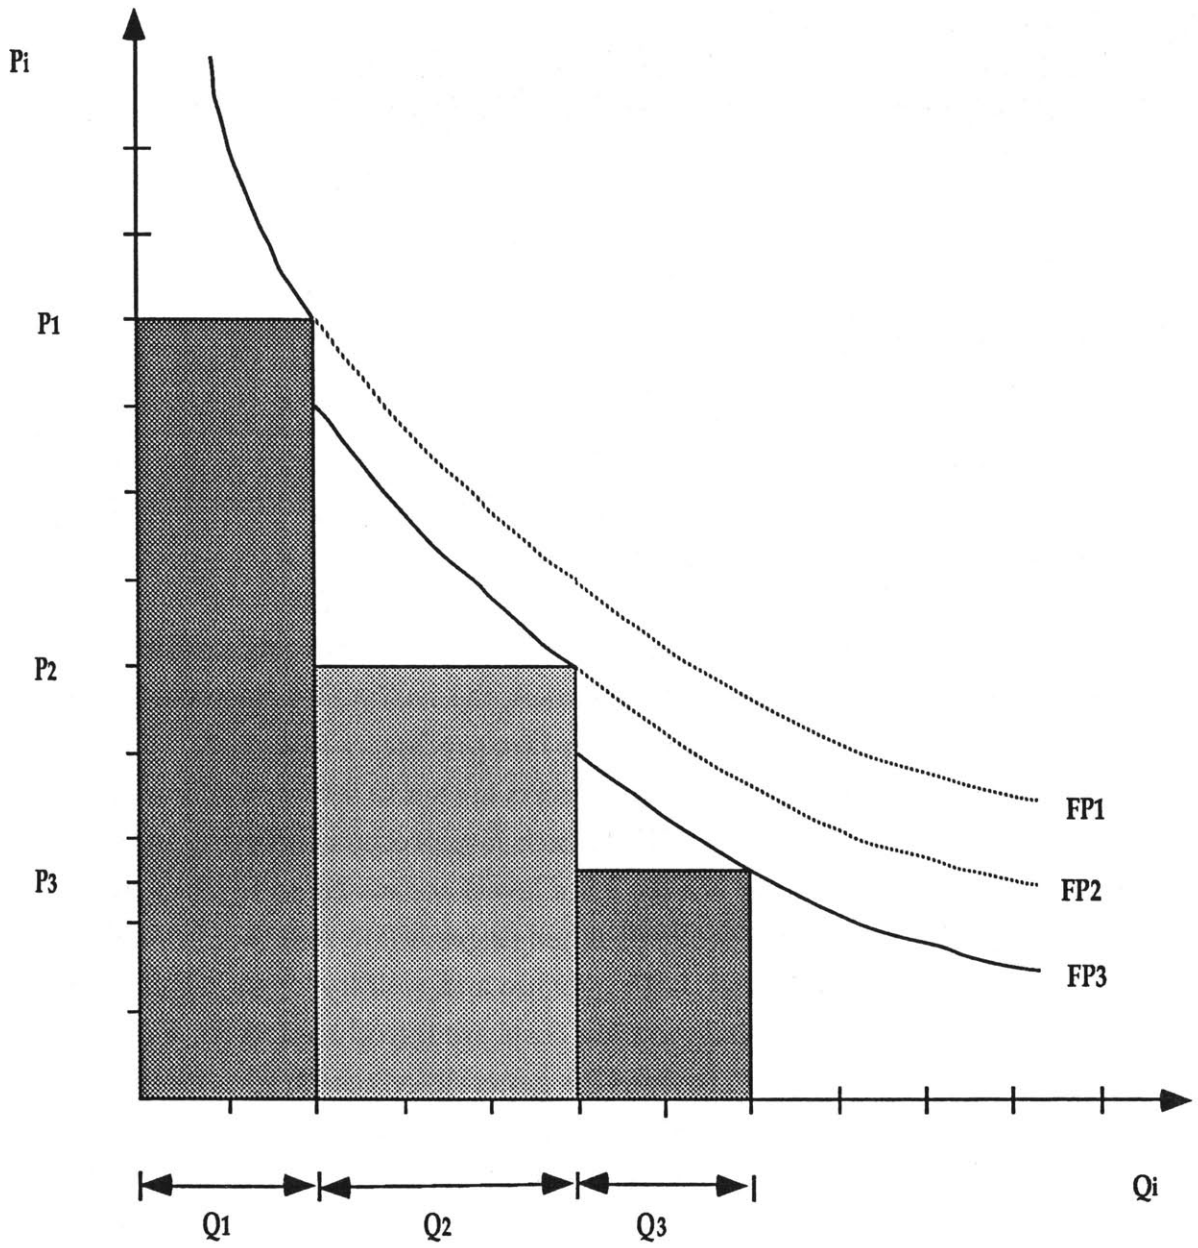
Figure 4.4 - Representation of OD Market Demand - Three Fare Product
Example

Generalized Cost Model
Three Fare Product Example
Consumer Surplus Calculation
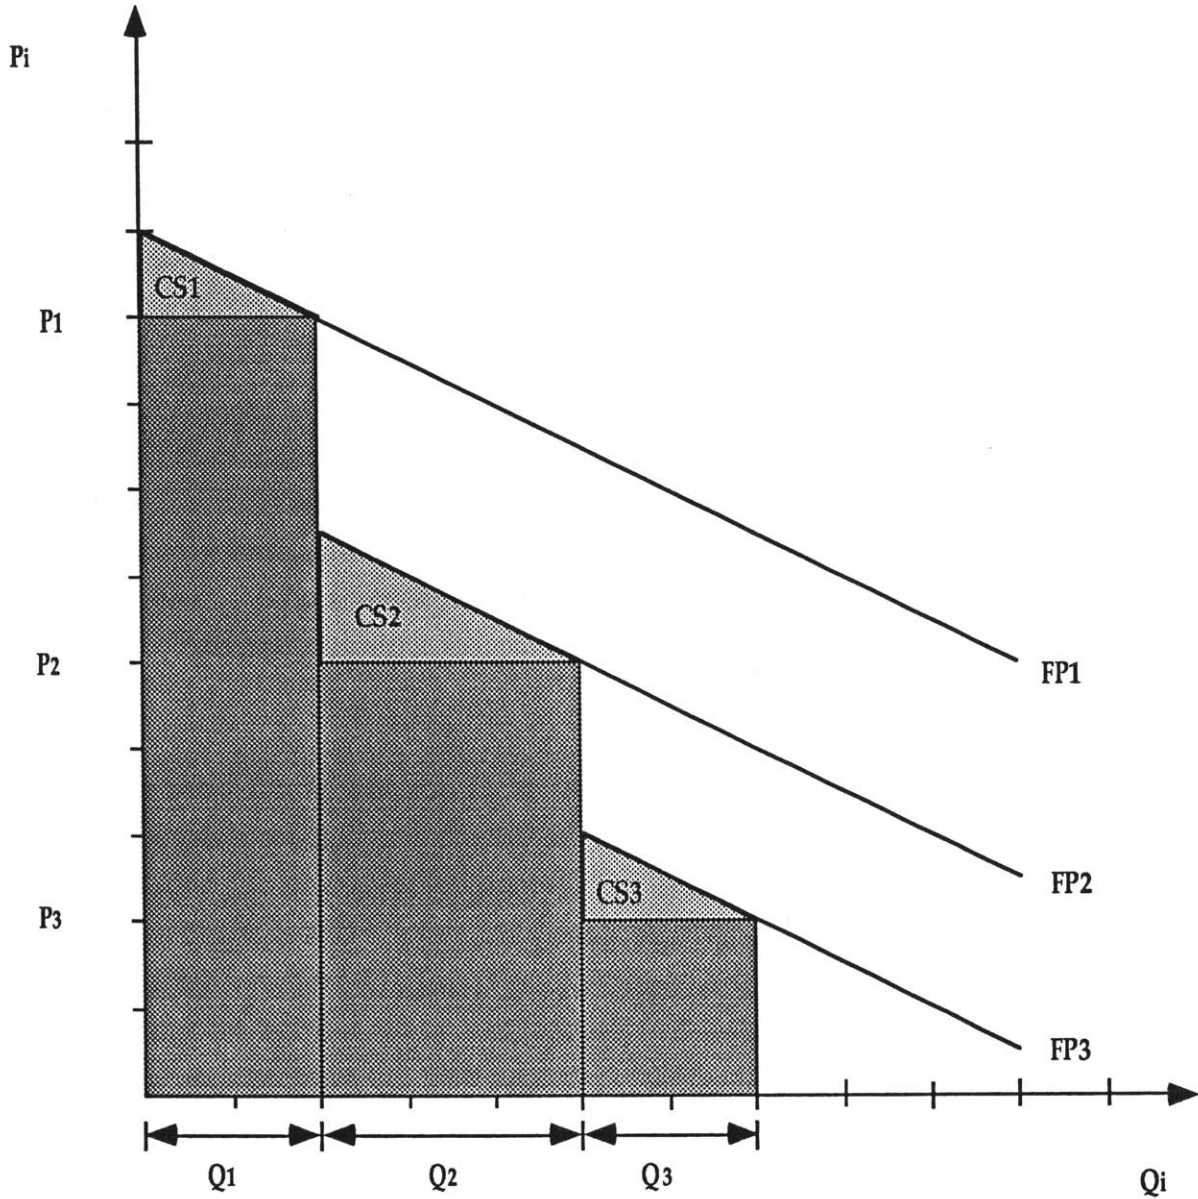
Figure 4.5 - Generalized Cost Model Three Fare Product Example - Consumer Surplus Calculation

range of passengers purchasing that fare product. Mathematically, the welfare losses resulting from the purchase restrictions accompanying fare products 2 and 3 are :

$$\text{Welfare Loss}_2 = \int_{Q_1}^{Q_1+Q_2} \left[f_1^1(q) - f_2^1(q) \right] dq \tag{4.8}$$

$$\text{Welfare Loss}_3 = \int_{Q_1+Q_2}^{Q_1+Q_2+Q_3} \left[f_1^1(q) - f_3^1(q) \right] dq \tag{4.9}$$

The welfare loss associated with a single fare product can be calculated since the composition of passengers by fare product is maintained in the example. Recall that the ordering of passengers by willingness to pay is maintained from fare product to fare product. It should also be noted that there is no deadweight welfare loss associated with fare product 1 since it is without purchase restrictions.

In the three fare product example with linear demands shown in Figure 4.6, the price levels P_1, P_{2zdc}, and P_{3zdc} represent the price levels that are charged to obtain the passenger mix Q_1, Q_2, and Q_3 in the model with zero degradation costs. The price levels P_1, P_2, and P_3 represent the price levels charged in the market to maintain the identical passenger mix (Q_1, Q_2, Q_3) when degradation costs are introduced to the model. The welfare loss to the market associated with the fare product 2 degradation costs is represented by the hatched area labeled WL22 in Figure 4.6. The loss associated with the fare product 3 degradation costs, on the other hand, is the sum of the hatched areas labeled WL32 and WL33 in the figure. The area WL32 represents the deadweight welfare loss associated with the purchase restriction bundle accompanying fare product 2 (as well as fare product 3) while the area WL33 represents the loss associated with the purchase restriction bundle accompanying only fare product 3.

Generalized Cost Model
Three Fare Product Example
Welfare Implications
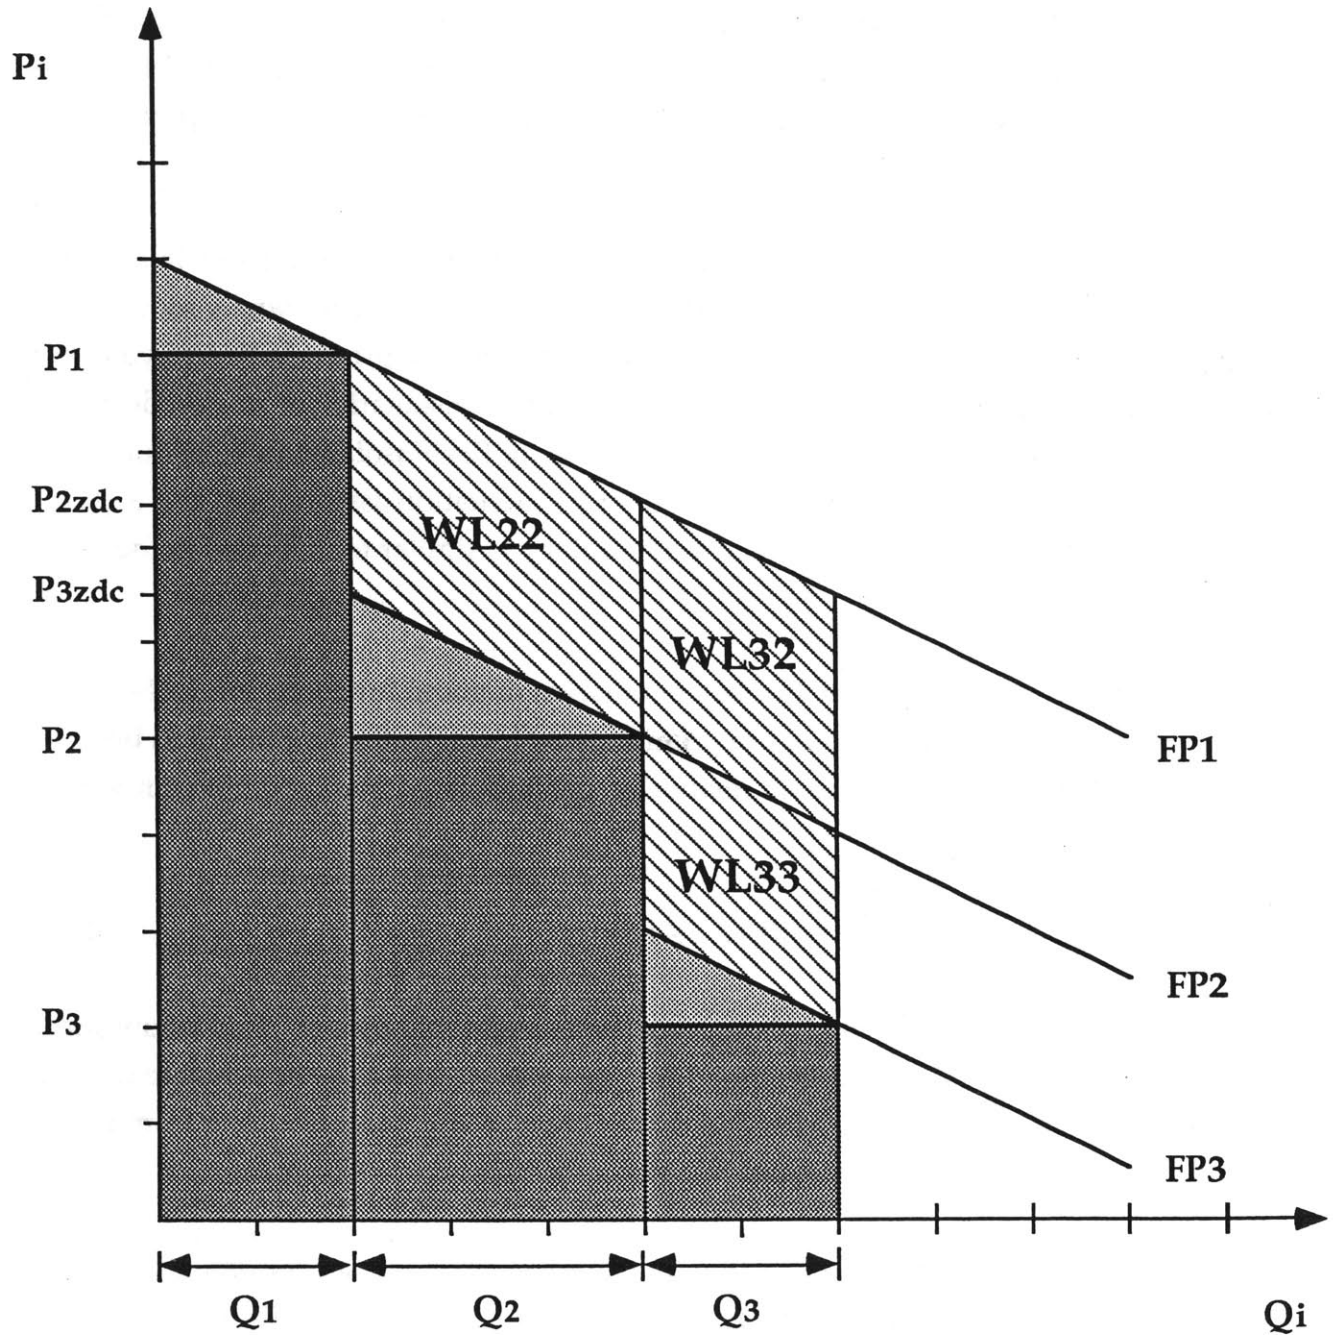
Figure 4.6 - Generalized Cost Model Three Fare Product Example - Welfare
Implications

4.3.2 General N Fare Product Model

Assessing the impacts of the fare product mix on traffic and revenue is contingent upon the ability to model all available fare products offered for sale on the market. Since airlines offer seemingly countless fare products in as many as fifteen different fare classes, the usefulness of the model depends upon the ability to extend it beyond the trivial cases. The simple three fare product model is extended to characterize the market demand in the general N fare product case in this section.

In the general case, N differentiated fare products are offered for sale in the OD market. The fare products must conform to the inferiority hierarchy motivated in Chapter One which requires that less-restricted fare products be equivalent or superior in all attributes except price level. A representation of market demand for the N fare product case is constructed using the N isolated fare product demand functions, the perfect passenger segmentation assumption, and a set of N fare product price levels. The methodology is similar to that used in the three fare product example.

As previously motivated, the isolated market demand functions for the N differentiated fare products are progressively lower as the fare products become increasingly restricted. Mathematically, the N isolated market demand functions must have the property :

$$f_N(P). \dots \leq f_{i+1}(P) \leq f_i(P) \dots \leq f_1(P) \tag{4.10}$$

for all values of P, where N is the most restricted fare product. This assumption ensures that the isolated demand functions used to build the final N fare product representation of market demand do not cross at any point resulting from non-rational consumer preferences.

Using the structure of the isolated demand functions, passenger demand can be allocated to the N fare products by assuming the ability of the airline to perfectly identify and segment the total demand population into N fare product subpopulations. Perfect segmentation assumes that passengers who are willing

to purchase a fare product are prevented from purchasing all more-restricted fare products as a result of their sensitivities to fare product attributes.

In the N fare product case, the passenger demand for any fare product i is given by the equation :

$$Q_i = f_i(P_i) - \sum_{j<i} Q_j \quad \forall\ i \in N \tag{4.11}$$

In other words, the number of passengers requesting fare product i is equal to the number of passengers that would have purchased fare product i at price level P_i had it been the only product on the market minus the number of passengers who are captive to one of the less-restricted fare products.

As before, a one to one correspondence must exist between the Q_i passengers captive to fare product i and the Q_i passengers with the related willingness to pay on the isolated market demand function for each increasingly restricted fare product (products i+1, . . . , N) to prevent double counting of individual passenger requests. Simply stated, the difference in the demand function must be assumed to maintain the order of willingness to pay at the level of the individual passenger. This requirement ensures that only the segmented passengers are removed from the more-restricted fare product demand functions. The remaining passengers are eligible to purchase a more-restricted fare product.

The N price levels will determine the shape of the representation of market demand as in the three fare product case. The price levels that are set for the N fare products must be strictly decreasing with increased restrictions. A case of simple dominance would result if a higher fare were charged for a more-restricted fare product. It is logical, assuming perfect information, that a dominated product will have no demand and can be ignored.

Setting the price levels P_i attracts Q_i passengers for each fare product i. The representation of market demand for the flight in the N fare product case can be constructed, given the price levels P_i for i = 1, . . . , N. The sum of all passengers using all fare products is :

135

$$Q_{total} = \sum_{i=1}^{N} Q_i \tag{4.12}$$

which gives the total demand for seats on the aircraft Q_{total}. A total revenue contribution to the airline of :

$$Rev_{total} = \sum_{i=1}^{N} P_i Q_i \tag{4.13}$$

is made by these Q_{total} passengers.

The consumer surplus associated with fare product i is calculated as :

$$Consumer\ Surplus_i = \int_{\sum_{j<i} Q_j}^{\sum_{j\leq i} Q_j} f_i^{-1}(q)dq - P_i Q_i \tag{4.14}$$

where each fare product j is less restricted than fare product i. The welfare associated with each fare product i is :

$$Welfare_i = \int_{\sum_{j<i} Q_j}^{\sum_{j\leq i} Q_j} f_i^{-1}(q)dq \tag{4.15}$$

while the deadweight welfare loss associated with each more-restricted fare product i equals :

$$Welfare\ Loss_i = \int_{\sum_{j<i} Q_j}^{\sum_{j\leq i} Q_j} \left[f_1^{-1}(q) - f_i^{-1}(q) \right] dq \tag{4.16}$$

136

where all variables and functions are defined as previously. Summing the consumer surplus associated with each fare product yields the total consumer surplus for the market. Market totals for welfare and welfare loss can be calculated analogously.

4.3.3 Review of the Basic Model Assumptions

This section has presented the framework for the model of airline fare product differentiation developed in this dissertation. Several assumptions have been made that should be made clear to the reader. To clarify the assumptions made in the model, they are listed below :

1) Single flight
2) Single carrier (no competition)
3) Single OD market
4) Single cabin of in-flight service
5) N differentiated fare products
6) Each fare product i+1 is more restricted than fare product i
7) Airlines develop fare products to <u>perfectly</u> identify and segment the population into subgroups by their willingness to pay
8) Subgroups are ordered by the willingness to pay of their members
9) A one to one correspondence exists between passengers on all isolated fare product demand functions
10) Invertible isolated demand functions
11) Price levels decrease with increased restrictions
12) Passengers incur costs associated with accepting additional purchase restrictions
13) Demand functions are always non-negative
14) Demand functions are non-increasing in P_i

The assumption that passengers incur costs associated with additional purchase restrictions (12) and the one to one correspondence existing between passengers on the invertible isolated demand functions (9 and 10) are the major differences compared to the assumptions of the standard monopoly price discrimination

model presented in Chapter Two. The balance of the assumptions are quite standard and straightforward.

4.4 Differences in Demand Functions : Incorporating Generalized Costs

Having motivated the framework for a static fare product differentiation model, it remains to make a version of the model operational. Since the motivated framework relies upon differing demand functions, the form of these differences must be quantified. But first, the major factors causing the demand function differences must be uncovered. A closer analysis of the reasons that air travel demand decreases for the more-restricted fare products will help to determine the nature of the demand differences.

The impacts of purchase restrictions on demand should be focused upon since, other than price level, the airline fare products considered here differ only by the level and severity of their restrictions. Two major effects that cause the demand functions for the more-restricted fare products to be lower relative to the less-restricted are 1) the possibility that the restrictions make the fare product unavailable to certain highly inflexible consumers and 2) the costs incurred by passengers associated with the inconvenience of the increased restrictions.

Availability effects have been assumed to have no effect on the calculation of the underlying demand function in the modeling framework. The possibility that certain fare products are unavailable to some consumers can be addressed alternatively by applying fences, capacity constraints, and booking limits to the underlying passenger demand functions. The universal fare product availability assumption is relaxed later in the chapter using the yield management techniques previously discussed. At present, however, all passengers are assumed to be able to purchase all fare products (although certain passengers have preferences among the fare products).

Unlike availability, the inconvenience associated with accepting an additional restriction is explicitly considered in the motivated passenger demand framework. The value that consumers have for air travel is expected to decrease as the level of restrictions on the fare products they seek to purchase increases.

The value decrease of the fare products is the primary cause for the demand functions to be lower for those consumers choosing to travel. Looked at another way, passengers incur a cost associated with the decreased convenience they experience due to the increased purchase restrictions facing them.

An example of potential costs incurred by air travelers associated with an increased level of purchase restrictions involves excursion-type tickets. When a fare product requires a Saturday night stay, the passenger is likely to incur a cost associated with spending an additional night in the same city. The cost may be thought of in terms of an extra night of hotel, day of rental car, and three additional meals, for example. More detailed measures might include the opportunity cost associated with a salesman remaining in a city an additional day versus traveling to another city with greater sales opportunities. Regardless of the factors considered, such costs drive the relative differences of the more-restricted fare product demand functions.

If each consumer is assumed to pay a cost associated with accepting an additional purchase restriction, the demand function for each increasingly restricted fare product i is lower according to a cost function, $c_i(\bullet)$, relative to fare product i-1. The function $c_i(\bullet)$ models the cost, referred to as the degradation cost, associated with accepting any purchase restrictions attached to fare product i that do not apply to fare product i-1. From the perspective of the consumer, it is as though the cost of the fare product has increased by the cost of meeting the increased purchase restrictions. For the airline, the cost incurred by passengers for accepting an additional restriction bundle translates into a lower-than-expected demand level for a given price level (had no cost been imposed on consumers). In other words, the passenger is paying a higher price for air travel but the airline is not earning any additional revenue as a result.

In order to construct an operational form of the motivated model, a functional form must be chosen for the cost function that both adheres to the assumptions of the modeling framework and accurately reflects the travel inconvenience experienced by passengers. The form of the cost function must not violate any of the assumptions of the model to ensure the validity of the motivated demand function. As with any demand function, the number of passengers requesting

service is constrained to be no less than zero at any price level. When incorporating the cost function, the demand function must satisfy the following :

$$Q_i = f_i(P_i, c_k(\cdot) \; \forall k \leq i) \geq 0 \qquad (4.17)$$

for all fare products i over the relevant ranges of P_i and c_i. In addition, the demand function $f_i(P_i, c_k(\bullet) \; \forall k \leq i)$ must be strictly non-increasing in P_i for all fare products $i \in N$ that are available in the OD market.

Finally, the ordering of passengers by their willingness to pay must be preserved for each more-restricted fare product demand function. In order for the willingness to pay hierarchy in the representation of market demand to remain intact, the ordering must be preserved at the level of the individual passenger. To ensure this, the cost differential between any two passengers may not exceed their differential in willingness to pay. This requirement will guarantee that, although the willingness to pay of passengers may change, no two passengers will change places in the willingness to pay hierarchy for the flight departure.

4.4.1 Generalized Cost Model - Joint Price Level Optimization Formulation

The generalized cost model can be used to calculate the optimal price levels for a given fare product structure. The basic form of the optimization appears below :

$$\text{Max } R = \sum_{i=1}^{N} P_i Q_i$$

Subject to :

$$\sum_{i=1}^{N} Q_i \leq Cap$$

$$Q_i \geq 0 \qquad \text{for } i = 1,...,N$$

140

where

R = total revenue

Q_i = number of seats allocated to fare product i

P_i = average fare charged to fare product i

N = total number of fare products

Cap = total aircraft capacity

c_k = cost associated with accepting the restriction bundle accompanying fare product k

The objective seeks to maximize total revenues for all fare products subject to an aircraft capacity constraint and demand level non-negativity constraints.

The generalized cost modeling framework requires that a passenger demand specification of the form :

$$Q_i = f(P_i, c_k(\cdot) \; \forall k \leq i) - \sum_{j < i} f(P_j, c_k(\cdot) \; \forall k \leq j) \qquad (4.18)$$

be assumed and substituted into the objective function. The joint price level optimization program can then be solved for the N fare product case using non-linear optimization techniques. The result would be the optimal price levels for each fare product and the resulting number of passengers expected to purchase them.

4.4.2 Limiting Cases of the Generalized Cost Model

The asymptotic limits provide some useful insight to the basic structure underlying the generalized cost model. Figure 4.7 shows the limit as the costs associated with accepting the purchase restrictions on the lower-priced fare products approaches zero in the three fare product case. If the costs associated with accepting the additional restrictions on all of the N fare products are zero, the model is identical to the standard monopoly price discrimination model described in Chapter Two. Thus, the standard monopoly price discrimination model is simply a special case of the generalized cost model.

A fare product structure having zero degradation costs for all fare products represents the best case scenario for the airlines since the demand functions do

Generalized Cost Model
With No Degradation Costs
Three Fare Product Example
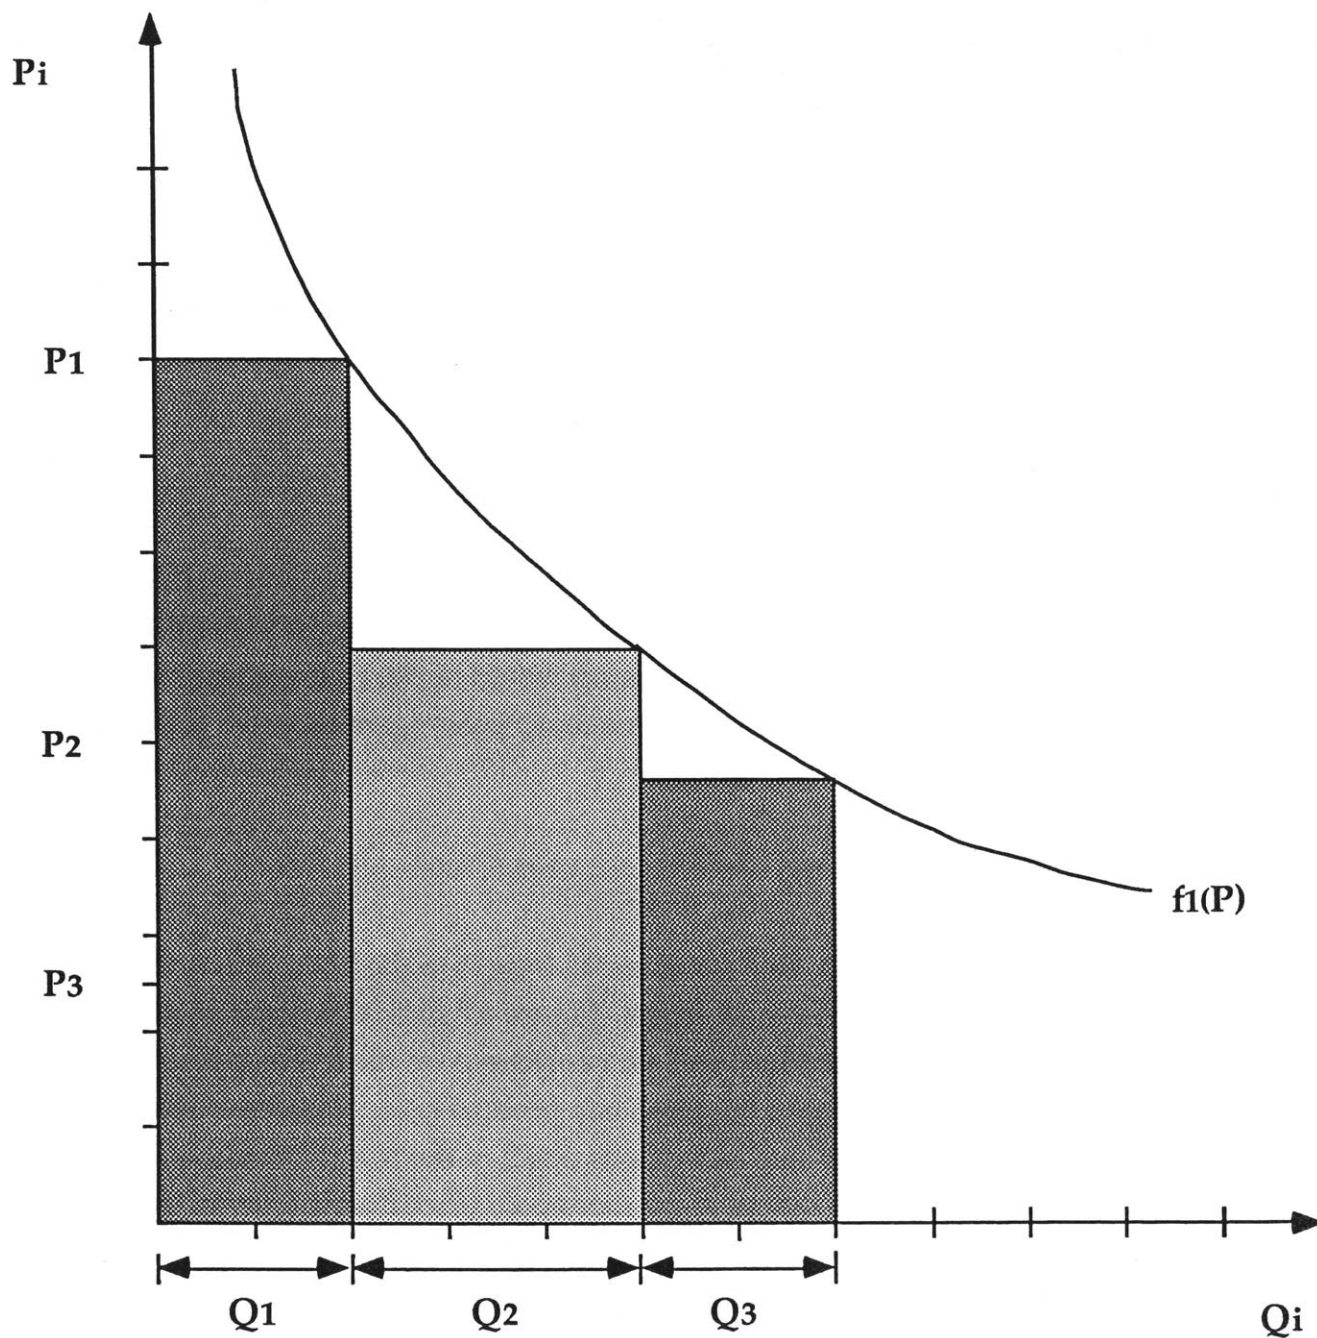
Figure 4.7 - Generalized Cost Model With No Degradation Costs - Three Fare
Product Example

not differ with the increased purchase restrictions. Airlines should strive to develop fare product restrictions that impose the least cost on consumers (when passenger segmentation devices are assumed to be perfect) to limit the magnitude of the demand function reductions of the restricted fare products. More passengers choose to travel at a given price level, or conversely, the same number of passengers travel at a higher fare level when degradation costs decrease.

As the cost associated with accepting the additional bundle of restrictions on a fare product becomes very large, on the other hand, the demand for the more-restricted fare products is driven to zero. At the extreme, the model becomes similar to a single product demand function as the costs associated with all but the unrestricted fare product are very large (approach infinity). This is intuitively pleasing since extreme costs to consumers are expected to drive the demand for any product to zero. Therefore, a single product market model is also a special case of the generalized cost model. It is interesting to note that when the number of fare products considered is one, then a single product demand function exists in the market as well. In any case, airlines should be aware that large perceived degradation costs prevent them from price discriminating in a way that maximizes revenue. Development of effective passenger segmentation devices (*e.g.* differentiated fare products) that impose realistic costs on consumers is required to earn incremental revenues through price discrimination even when optimizing price levels.

4.4.3 Pareto Optimality and the Generalized Cost Model

In the limiting case where the degradation costs are equal to zero for each of the more-restricted fare products, Pareto optimality can exist. A Pareto optimal situation occurs when no consumer can be made better off by employing a different product allocation process. To achieve a Pareto optimal solution in the case of fare product differentiation, the price level of the unrestricted fare product must not exceed its price level had no other fare product been offered on the market. Although Pareto optimality is by no means guaranteed in the case of zero degradation costs, it is possible.

In contrast, Pareto optimality cannot exist when passengers incur degradation costs as a result of the imposed purchase restrictions. The deadweight welfare loss that occurs from passengers purchasing restricted fare products for which they experience degradation costs prevents the allocation of the available seats from being Pareto optimal. This is clear since an airline seat allocation in which no degradation costs are incurred would unambiguously benefit all consumers purchasing the restricted fare products.

A limiting case of the generalized cost model can result in a Pareto optimal situation. As degradation costs approach infinity, demand for the restricted fare products is driven to zero. In this case, society does not bear the deadweight welfare loss resulting from degradation costs since no passenger purchases any of the restricted fare products. The generalized cost model with infinite degradation costs mirrors the single, unrestricted fare product case and Pareto optimality exists for the market when only a single fare product is offered since no passenger can be made better off by a change in the allocation process. Even one passenger purchasing a restricted fare product with a non-zero degradation cost, however, would result in a loss of the Pareto optimal solution.

It is unlikely that Pareto optimality exists under the current practice of airline fare product differentiation. Lower-priced, more-restricted fare products are sold to passengers and passengers do incur costs associated with accepting the attached purchase restrictions. Since many passengers incur degradation costs as the result of the purchase restrictions introduced into the allocation process by the fare product differentiating airlines, societal welfare is reduced and thus, the Pareto optimal solution is lost. Simply stated, price discrimination techniques that impose costs upon consumers in order to segment them can never result in a Pareto optimal situation.

Previous works have hypothesized that Pareto optimality can exist with airline fare product differentiation (e.g. Belobaba, 1987) based upon analysis using the standard monopoly price discrimination model as a market characterization. The standard monopoly price discrimination model does not, however, take into account the costs incurred by passengers as a result of the imposed purchase restrictions. The generalized cost model takes into account these degradation costs which, in turn, prevent Pareto optimality from being achieved. Thus, the

144

only versions of the generalized cost model under which Pareto optimality can be achieved are those that impose no degradation costs on the consumer population.

From a Pareto optimality standpoint, airline fare product differentiation is not completely justified unless a segmentation technique which imposes no degradation costs on consumers can be developed. The result of its application would, of course, need to be shown to be Pareto optimal. Airlines do, however, have incentive to develop fare products which minimize degradation costs since any improvement towards achieving Pareto optimality translates directly into increased revenue. The revenue incentive provides the mechanism for change and, potentially, the achievement of a Pareto optimal situation through improvements in the fare product structure. The reader is reminded that Pareto optimality is not the sole measure of efficiency that benefits society. The benefits provided by allocative efficiency (as discussed in Chapter Two of this dissertation) are achievable under the current structure of airline fare product differentiation and revenue management.

4.5 Specifying a Form for the Generalized Costs

To make the generalized cost model operational, a specific functional form must be assumed for the degradation costs facing consumers. In this section, three different functional forms are proposed. The three models presented incorporate the costs associated with consumer inconvenience while satisfying the assumptions of the motivated framework. Each form can be used to characterize the impact of purchase restrictions on passenger demand making different assumptions about the nature of the degradation costs facing consumers. Not only does the generalized cost model contribute the first explicit consideration of the costs incurred by passengers that result from applied purchase restrictions to appear in the airline yield management literature, it also allows for the flexible specification of those degradation costs.

4.5.1 Constant Cost Model

Many of the out-of-pocket costs incurred by consumers faced with an additional purchase restriction are similar. For instance, the cost associated with an extra night in a hotel, an extra day of the rental car, and three additional meals can be assumed approximately equal for all consumers. The constant cost model assumes that each consumer pays an identical and constant fixed cost associated with accepting each additional purchase restriction. In other words, the form of the cost function (or demand difference) for each of the isolated restricted fare product demand functions is a constant. The constant cost, of course, can be varied to represent different levels of consumer inconvenience.

The relationships between the isolated restricted fare product demand functions can be specified using an assumption about cost and the fare product inferiority hierarchy. In the constant cost model, the demand function for a more-restricted fare product is lower by a constant amount for any given price level when compared to the next-less-restricted fare product. In other words, the price that passengers are willing to pay for the more-restricted fare product i is reduced by a constant, c_i, when compared to the less-restricted fare product i-1 :

$$f_i(P_i) = f_{i-1} (P_i + c_i) \qquad (4.19)$$

where c_i is the cost to all consumers of accepting the additional restrictions placed on fare product i not found in the less-restricted fare product i-1.

Returning to the three fare product example, the relative positions of the isolated demand functions are presented in Figure 4.8. In the figure, the demand function for fare product 2 is lower by the value c_2 when compared to the unrestricted product while the demand function for fare product 3 is lower by c_3 relative to fare product 2. The degradation cost c_2 would be incurred by passengers as a result of being required to meet the 3 day advance purchase restriction if, for example, fare product 2 were an AP3 fare product. The cost c_3 would measure the added burden to passengers of extending the advance purchase requirement an extra 4 days, requiring a round trip purchase, and imposing a Saturday night stay at the destination if fare product 3 were an E7NR fare product.

146

Constant Cost Model for
More-Restricted Airline Fare Products
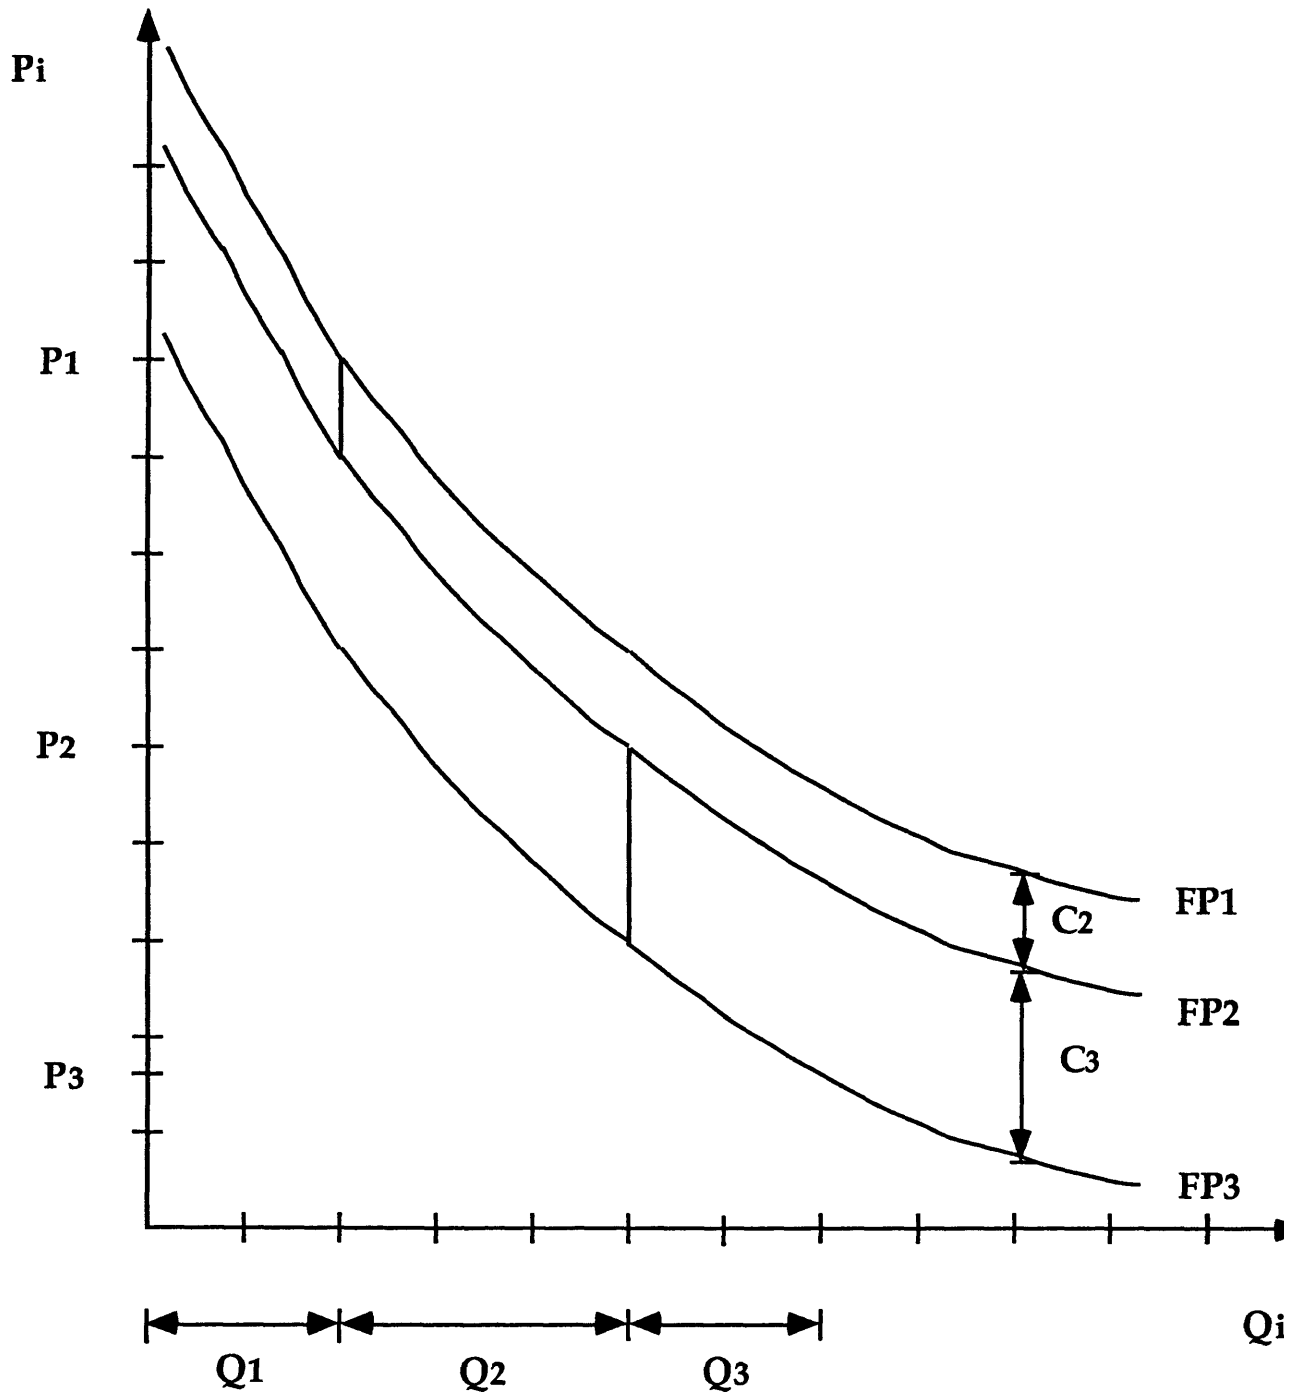
Figure 4.8 - Constant Cost Model for More-Restricted Fare Products

147

In the isolated demand functions for the more-restricted fare products, it is clear that the ordering of individual passengers by their willingness to pay is not altered since the cost incurred by each individual consumer is identical. Simple mathematics tells us that the subtraction of a constant from all members of a strictly ordered set (of integers) will not change the rankings of its members. Thus, the identical passenger ordering assumption required in the modeling framework is preserved in the constant cost model.

Once the constant cost associated with accepting each additional purchase restriction has been identified, the demand functions for all fare products can be determined from any single fare product demand function. Since all passengers receive the same in-flight amenities and the fare products differ only in the applied purchase restrictions, it is convenient to think of the more-restricted fare products as being derived from the unrestricted fare product. In this framework, the demand function for the unrestricted fare product provides an upper bound on the number of passengers wishing to travel on the single flight offered in the isolated OD market at any price level P because the unrestricted product represents the best possible service option.

The only two sets of information required to calculate the isolated demand functions for all N airline fare products are 1) the unrestricted fare product demand function and 2) the N costs imposed by each additional purchase restriction bundle. Since no cost is incurred by consumers as a result of purchasing the unrestricted fare product, the value c_1 is set to zero. In terms of the unrestricted fare product, the isolated demand function for any other fare product i can be calculated as :

$$f_i(P_i) = f_1(P_i + t_i) \tag{4.20}$$

where t_i represents the total degradation cost associated with fare product i relative to the unrestricted fare product. It should be noted that if the entire fare product structure were in place, the total degradation cost associated with fare product i, t_i, would be equal to $\sum_{j \leq i} c_j$. As before, $f_1(\bullet)$ represents the isolated unrestricted fare product demand function. A specific three fare product example demonstrates the relation between the more-restricted products and the

148

unrestricted fare product with Figure 4.8 showing that fare product 2 is lower by $t_2 = c_2$ compared to the unrestricted product while fare product 3 is lower by $t_3 = c_2 + c_3$.

Under the assumption of perfect passenger identification and segmentation, setting the price levels for the N differentiated fare products enables the calculation of the demand on the flight departure. The number of passengers purchasing fare product i at any given price level P_i is calculated according to equation 4.11. Substituting the unrestricted fare product demand function and the costs incurred by consumers (equation 4.20) into the general modeling framework passenger demand equation (equation 4.11) yields the passenger demand for any fare product i in terms of the unrestricted fare product demand function and the costs of the restriction bundles :

$$Q_i = f_1(P_i + t_i) - \sum_{j<i} Q_j \quad \forall\ i \in N \tag{4.21}$$

where each j represents a fare product that is less restricted than fare product i and t_i represents the total degradation cost associated with fare product i relative to the unrestricted fare product. The demand for the N differentiated fare products can now be calculated along with the revenue generated.

4.5.2 Increasing Cost Model

The constant cost model is by no means the only specification of perceived airline fare product degradation costs that can be incorporated into the generalized cost framework. In certain situations, the cost incurred may differ by passenger rather than being constant. For instance, while the cost incurred for an extra day of rental car and hotel is well represented as a fixed and constant cost, the opportunity cost associated with spending a Saturday night away would be expected to vary with the value of time of a particular individual. Since the value of time is expected to vary greatly between individuals, a variable cost would seem to be a more appropriate representation of the opportunity cost incurred by an individual. In cases where the opportunity costs associated with a bundle of restrictions dominate the out-of-pocket costs, a variable cost specification may be preferable.

149

It would seem logical that larger value of time costs would be positively correlated with higher values of willingness to pay. Certainly, travel itinerary flexibility has more importance to individuals who place greater value on their time. Passengers having high values of willingness to pay are expected to be affected to a greater extent by the restrictions that seek to limit their travel itinerary flexibility. In this case, it may be appropriate to model the fare product degradation costs facing consumers as an increasing function of willingness to pay. In the generalized cost framework, a model can be constructed by assuming that the cost associated with accepting increased purchase restrictions is an increasing function of willingness to pay.

The increasing cost model described in this section allows for the calculation of the passenger demand function for the N fare products offered in the market with increasing costs to passengers depending upon their willingness to pay for air travel. Mathematically, the cost functions $c_i(Q_i)$ must be non-increasing in Q_i for all fare products $i \in N$ where Q_i represents the vector of passenger demands $Q_i = (Q_j : j = 1, \ldots, i)$. The preservation of the motivated demand function required for the generalized cost model relies upon the form of costs facing consumers not disrupting the ordering of consumers by their willingness to pay. As before, the ordering of the individual passengers by their values of willingness to pay must remain the same across all restricted fare product demand functions. Thus, any non-increasing cost function $c_i(Q_i)$ can be used in the model, provided that no two passengers switch places on the demand function $f_i(P_i)$. Unfortunately, the passenger willingness to pay ordering is not guaranteed for all increasing cost functions as it was in the constant cost model. Thus, the cost function must be checked to determine if it satisfies the ordinality condition for consumer willingness to pay.

Finally, the computational burden associated with the increasing cost model is expected to be greater than that of the constant cost model because a function associated with the incremental cost of each restricted fare product must be obtained instead of a simple constant. The ability to obtain a serviceable cost function that correctly represents the experience of passengers must be weighed against the benefits of using a variable cost model.

4.5.3 Constant and Increasing Cost Model

While compelling cases can be made for both the constant and increasing cost models as the most appropriate form of consumer costs, a model incorporating both types of costs is better still. The constant and increasing cost model incorporates both an increasing cost component and a constant cost component in the analysis. The two term degradation cost specification provides increased modeling flexibility.

The cost function $C_i(Q_i) = k_i + c_i(Q_i)$ presented in this version of the model accounts for the constant cost incurred by passengers in the constant term k_i, as well as the increasing cost captured by the function $c_i(Q_i)$. Together, the two cost terms capture both the fixed costs associated with restrictions that are incurred by all passengers as well as the costs that differ by individual. The assumption of passengers having higher values of willingness to pay incurring greater costs remains for the variable cost component and, thus, the cost function overall. As in the previous model specifications, passengers may not switch places on the demand function in terms of willingness to pay. The requirements to preserve ordinality are analogous to those of the increasing cost model since the two models differ only by the inclusion of the constant cost term which has no effect on order.

Unfortunately, the calculation of the cost functions associated with the different fare products becomes increasingly difficult with the additional flexibility in the framework. The increased data requirement associated with the more flexible formulation may prove prohibitive in practice.

4.6 Modeling Passenger Diversion

To this point, the ability of the airline to identify and segment passengers by their willingness to pay has been implicit in the model. Although airlines would like to believe that such a segmentation is possible, consumers and airlines alike know that passengers can avoid even the most elaborate segmentation schemes. In the airline industry, passenger buy down behavior has been called <u>diversion</u>.

151

This section addresses the ability of passengers to avoid airline segmentation by incorporating diversion into the generalized cost model.

The section begins with a motivation of passenger diversion behavior. Next, the treatment of passenger diversion in the airline yield management literature is reviewed. Passenger diversion is then incorporated into the generalized cost model of airline fare product differentiation for the general N fare product case. The passenger diversion measure is extended to allow for the diversion of fare product passengers as a function of factors such as fare product price differentials. Finally, the limiting cases of passenger diversion are discussed.

4.6.1 Passenger Diversion Motivation

The nature of airline fare product differentiation must be examined to correctly motivate passenger diversion to the lower-priced fare products. In the current airline environment, almost everyone who has considered the purchase of air travel has been offered an array of fare products from which to choose. The many fare products have resulted from the use of second degree price discrimination or "self-selection" techniques by airlines to identify and segment the consumer population on the basis of their willingness to pay.

Airlines attempt to segment passengers through a fare product design scheme that is based on the belief that willingness to pay and the ability to meet purchase restrictions are highly correlated. The purchase restrictions accompanying the fare products are intended to make passengers self-select into mutually exclusive market segments based on their sensitivity to travel itinerary flexibility. It is the hope of airlines that all consumers choose the most expensive fare products within their travel budget as a result of the purchase restrictions making the less expensive options unattractive to them.

Unfortunately for the airlines, the self-selection techniques that they employ are quite fallible. With passenger diversion, as is often true of imperfections in market segmentation attempts, surplus is returned to the consumer at the expense of airline revenue. Passengers commonly purchase air travel services at prices significantly below what they are willing to pay. Price-insensitive passengers frequently can meet the restrictions accompanying the lower-priced

fare products and, thus, are able to buy down. For instance, if a business traveler knows more than a week in advance of a given flight departure that his travel plans are firm, he can purchase a three or seven day advance purchase fare product in lieu of the unrestricted fare product. The discount purchase occurs even though the traveler may have been willing to purchase the unrestricted product in the absence of an alternative.

The strict definition of passenger diversion in this dissertation refers to a consumer purchasing a lower-priced fare product than is intended by the airline in its attempt to price discriminate. To have incentive to divert, the perceived total value of the lower-priced fare product must be greater than that of the higher-priced. Passengers have incentive to divert only if the intrinsic value (e.g. not including price level) that they place on the higher-priced fare product exceeds that of the lower-priced fare product by an amount less than the price differential of the two products. Clearly, diversion would not occur if the consumers valued the higher-priced fare product more than the lower-priced. Thus, among all diverting passengers, the lower-priced product must be preferred to the higher-priced. Moreover, passenger diversion need not be dependent upon the application of booking limits to the fare products competing for the available aircraft capacity. The only requirement for passenger diversion is the existence of differentiated fare products that attempt to segment the population. Viewed in this way, passenger diversion is simply an avoidance of airline segmentation techniques.

Airlines could attempt to prevent passenger diversion by modifying their fare product attributes to better segment the population. To prevent passenger and travel agent confusion, however, the restrictions associated with the market selling fares must remain relatively stable over time. Airlines apply booking limits to the lower-priced fare products as an alternative to frequent restriction changes in the attempt to limit the revenue dilution occurring from diversion. Booking limits may induce passengers to purchase higher-priced fare products when discount services are unavailable on their preferred travel itineraries. Passengers are said to "sell up" when they purchase a higher-priced fare product after being denied a discount fare product. The downside to using booking limits to induce sell up is that some passengers will travel on a competing carrier or choose not to travel at all when denied their preferred fare product or flight

itinerary. It is interesting to note that the only passengers who sell up are those willing to pay enough to purchase the higher-priced fare products. In other words, only the diverting passengers are candidates to sell up.

To date, the shifting of passenger demand between the fare products has not been thoroughly addressed in the airline yield management literature. The modeling of passenger switching behavior between fare products has been limited to the sell up resulting from denied requests (to the exclusion of true diversion behavior). The next section reviews the literature related to passenger switching between fare products and motivates the need for more research.

4.6.2 Passenger Diversion Literature Review

The topic of passenger diversion to lower-priced fare products has received limited attention in the airline yield management literature. A survey article by Weatherford and Bodily provided a review of the airline yield management literature that sought to incorporate passenger diversion (Weatherford and Bodily, 1992). The article indicated only three relevant references (Belobaba, 1987, Pfeifer, 1989, and Brumelle et al., 1990). The term passenger diversion is used quite loosely in the survey article to refer to any relaxation of the assumption that passengers are divided into distinct and independent demand populations. The lack of references provided with such a broad definition of passenger diversion makes the gap in the literature concerning passenger buy down behavior apparent.

In his doctoral dissertation, Belobaba tested the effects of passenger sell up to higher fare classes on the expected marginal seat revenue (EMSR) seat allocation optimization for the two fare class case (Belobaba, 1987). Belobaba defined the probability of a passenger making a vertical shift from fare class i to fare class i-1 on the same flight leg as $P_i(v)$. The shift probability was incorporated into the EMSR seat allocation formulation and used to calculate an incremental seat protection level having accounted for expected passenger sell up.

The vertical choice shift probability required that a passenger be refused a seat for the discount fare product prior to considering the fare class upgrade. The measure used did not consider the ability of consumers to meet the purchase

restrictions imposed by the discount fare products, but rather, automatically assumed that they preferred but were unable to purchase the discount product. Thus, the result of Belobaba was actually an exploration of passenger sell up behavior and not diversion (as defined here) that required an unrealistic behavioral assumption.

Pfeifer modeled what he called passenger diversion to provide an optimal decision rule for calculating the discount fare class seat allocation in the two fare class case (Pfeifer, 1989). He assumed two types of passengers within the global fare product population, shoppers and non-shoppers. Shoppers would not purchase the higher-priced fare product when denied the discount while non-shoppers would. Again, all passengers were assumed to prefer the discount fare product unambiguously if it were available. Pfeifer's method also assumed the ability of all consumers to meet the requirements of the discount fare product purchase restrictions. Although the effects of passenger diversion can be measured using the decision rule developed by Pfeifer, the assumption of discount fare product dominance does not accurately reflect consumer behavior, especially when considering the effects of second degree price discrimination techniques used by airlines.

Brumelle *et al.* modeled passenger switching behavior as the probability that passengers will sell up to a higher booking class when the lower booking class is closed for reservations (Brumelle *et al.*, 1990). The result of their investigation was an optimal discount seat allocation subject to an upgrade probability attached to the otherwise independent demand populations (similar to Pfeifer). Again, the methods employed by Brumelle *et al.* presented what is more accurately referred to as the passenger sell up probability which is dependent upon booking limits initially denying service to passengers at the discount fare.

A rigorous behavioral motivation of passenger switching behavior has not yet been presented in the literature. The existing work has focused upon shifting percentages to higher-priced fare products without motivating the underlying behavior of consumers within the framework of the differentiated fare products. Contrary to the treatment received in the literature, passenger switching behavior is not exclusively the result of booking limits and capacity controls inducing the purchase of higher-priced fare products. Passenger switching more

commonly results from the ability of passengers to avoid the segmentation attempts of airlines. Since passengers must initially divert in order to be candidates to sell up, diversion is, by definition, more prevalent than sell up.

Strictly speaking, passenger buy down to the lower-priced fare products resulting from the ability to meet the imposed fare product purchase restrictions has not been considered in any of the modeling approaches claiming to address diversion. Rather, only the probability that a passenger will purchase the higher-priced fare product contingent upon being denied the discount fare product, or passenger sell up, has been addressed. True diversion, or passenger buy down, under airline fare product differentiation remains to be modeled.

Finally, the treatment of passenger switching between the different fare products was limited to the two fare class case in all previous articles. There remains a need for addressing the issues of passenger diversion to the lower-priced fare products and sell up for the general case of N fare products. Both passenger diversion as well as sell up can be modeled by extending the generalized cost model presented in this research. The diversion modeling methodology appears in the next subsection for the general N fare product case. The discussion of passenger sell up as it relates to the generalized cost model is postponed until the following section when booking limits are introduced. The modeling of true passenger diversion and sell up are performed for the first time in this dissertation for the general N fare product case, thus extending the airline yield management literature.

4.6.3 Modeling Passenger Diversion in the Generalized Cost Framework

Without the incorporation of passenger diversion, the generalized cost model contains several unrealistic assumptions. For instance, it assumes that airlines have the ability to perfectly identify and segment the population into distinct groups by their values of willingness to pay. The passenger identification and segmentation depends upon the price levels of the fare product in this framework. Thus, not only would the screening device be required to perfectly identify and segment the population, but also must be able to do so at any prevailing set of market price levels.

In addition, passengers falling within the arbitrarily set ranges of willingness to pay are assumed to unambiguously prefer their designated fare products. Passengers may only purchase the fare products that the airline intends according to their willingness to pay and the segmentation scheme. The self-selection devices used by airlines are, however, far from flawless, leaving a perfect identification and segmentation unachievable given currently employed price discrimination techniques. Therefore, realistic modeling requires that passenger buy down behavior be addressed explicitly.

<u>Mathematical Incorporation of Diversion into the Model</u>

The incorporation of passenger diversion into the generalized cost model framework requires changes to the objective function of the basic model formulation. Instead of modeling the revenue received by the airline as the number of passengers assumed to be segmented into purchasing the fare product intended, the objective function must allow for the revenue dilution experienced by passengers able to avoid the airline segmentation scheme. The exact change in the objective function depends upon the assumed form of passenger diversion. The calculation of the number of diverting passengers can be performed in one of two ways: as a percentage of the expected number of passengers or as an absolute number of passengers. The number of diverting passengers can be a fixed number or percentage and may be calculated as a function of either exogenous or endogenous modeling factors.

The first case considered assumes that the number diverting is a fixed percentage of the passengers expected in any fare product population at the prevailing price levels. Intuitively, the number of passengers able to divert to the lower-priced fare products should be related, in some way, to the number of consumers expected to purchase the higher-priced fare products. Calculating diversion as a percentage of the expected passengers satisfies this intuition. The model of diversion can be extended to make passenger diversion percentages a function of exogenous factors calculated *a priori* or endogenous factors to be included in the joint price level optimization formulation.

A concise objective function calculating passenger diversion as an absolute number would be difficult to implement since the number of passengers

157

diverting must be less than the total population for the intended fare product. This, of course, is guaranteed using the percentage of the population formulation previously proposed. In the case of an absolute number of diverting passengers, however, additional constraints would be required in the formulation to prevent more diversion than is possible from the actual fare product population size. The operational difficulties associated with incorporating an absolute number of diverting passengers remove it from consideration here.

The objective function of the model for the N fare product case incorporating passenger diversion to the lower-priced fare products as a percentage of passengers expected to purchase the higher-priced fare products appears below :

$$\text{Max } R = \sum_{i=1}^{N} P_i \cdot (1 - \sum_{j=i+1}^{N} d_{ij}) Q_i + \sum_{i=1}^{N-1} \sum_{j=i+1}^{N} P_j d_{ij} Q_i$$

where :

d_{ij} = % of fare product i passengers who divert to lower-priced fare product j

The revenue received by the airline is equal to the revenue expected from each fare product (without consideration of diversion) minus the decreased revenue associated with the loss of passengers to the lower-priced fare products plus the amount of revenue gained in the lower-priced fare products from the diverting passengers. The constraints imposed on the revenue optimization are identical to the case without diversion, as presented in the basic model description.

The single function representing the demand for OD market air travel for each isolated fare product reflects the maximum willingness to pay of each consumer (as in any demand function). Thus, the initial segmentation performed by the airline assumes that passengers will purchase the highest-priced fare product that they can afford. For this reason, passengers are not willing to purchase (e.g. do not divert to) higher-priced fare products, leaving passenger diversion to occur only to the more-restricted fare products.

158

In the diversion formulation, the variables P_i represent the optimal price levels that <u>should</u> be set to maximize revenues. The variables Q_i represent the mapping of the price levels and degradation costs from the demand functions of the form :

$$Q_i = f(P_i, c_k(\cdot) \; \forall k \leq i) - \sum_{j < i} f(P_j, c_k(\cdot) \; \forall k \leq j) \tag{4.22}$$

The variables Q_i, however, no longer represent the number of passenger expected to arrive in each of the fare product groups but rather the number of passengers that would book each fare product i at price level P_i provided that diversion did not exist. A new measure is required to represent the actual number of passengers expected to book each fare product i. For each fare product i, the number of passengers expected, q_i, is defined as :

$$q_i = (1 - \sum_{j=i+1}^{N} d_{ij})Q_i + \sum_{j=1}^{i-1} d_{ji}Q_j \tag{4.23}$$

where the variables are defined as before. The number of passengers expected to purchase fare product i is simply the number of passengers expected without diversion, Q_i, minus the number of fare product i passengers who divert to lower-priced fare products from i, $\sum_{j=i+1}^{N} d_{ij}Q_i$, plus the number of passengers diverting to fare product i from the more expensive fare products j.

The increase in consumer surplus associated with passenger diversion between less-restricted fare product i and more-restricted fare product j is equal to :

$$\text{Consumer Surplus Increase}_{ij} = \left[P_i - \sum_{i < k \leq j} c_k - P_j \right] d_{ij}Q_i \tag{4.24}$$

with the variables and functions as previously defined. Figure 4.9 shows the change in consumer surplus when diversion occurs exclusively from fare product 1 to fare product 2 in the three fare product example. The hatched box in the figure represents the consumer surplus that is gained by the passengers diverting from the unrestricted fare product to fare product 2.

159

Generalized Cost Model With Diversion
Three Fare Product Example
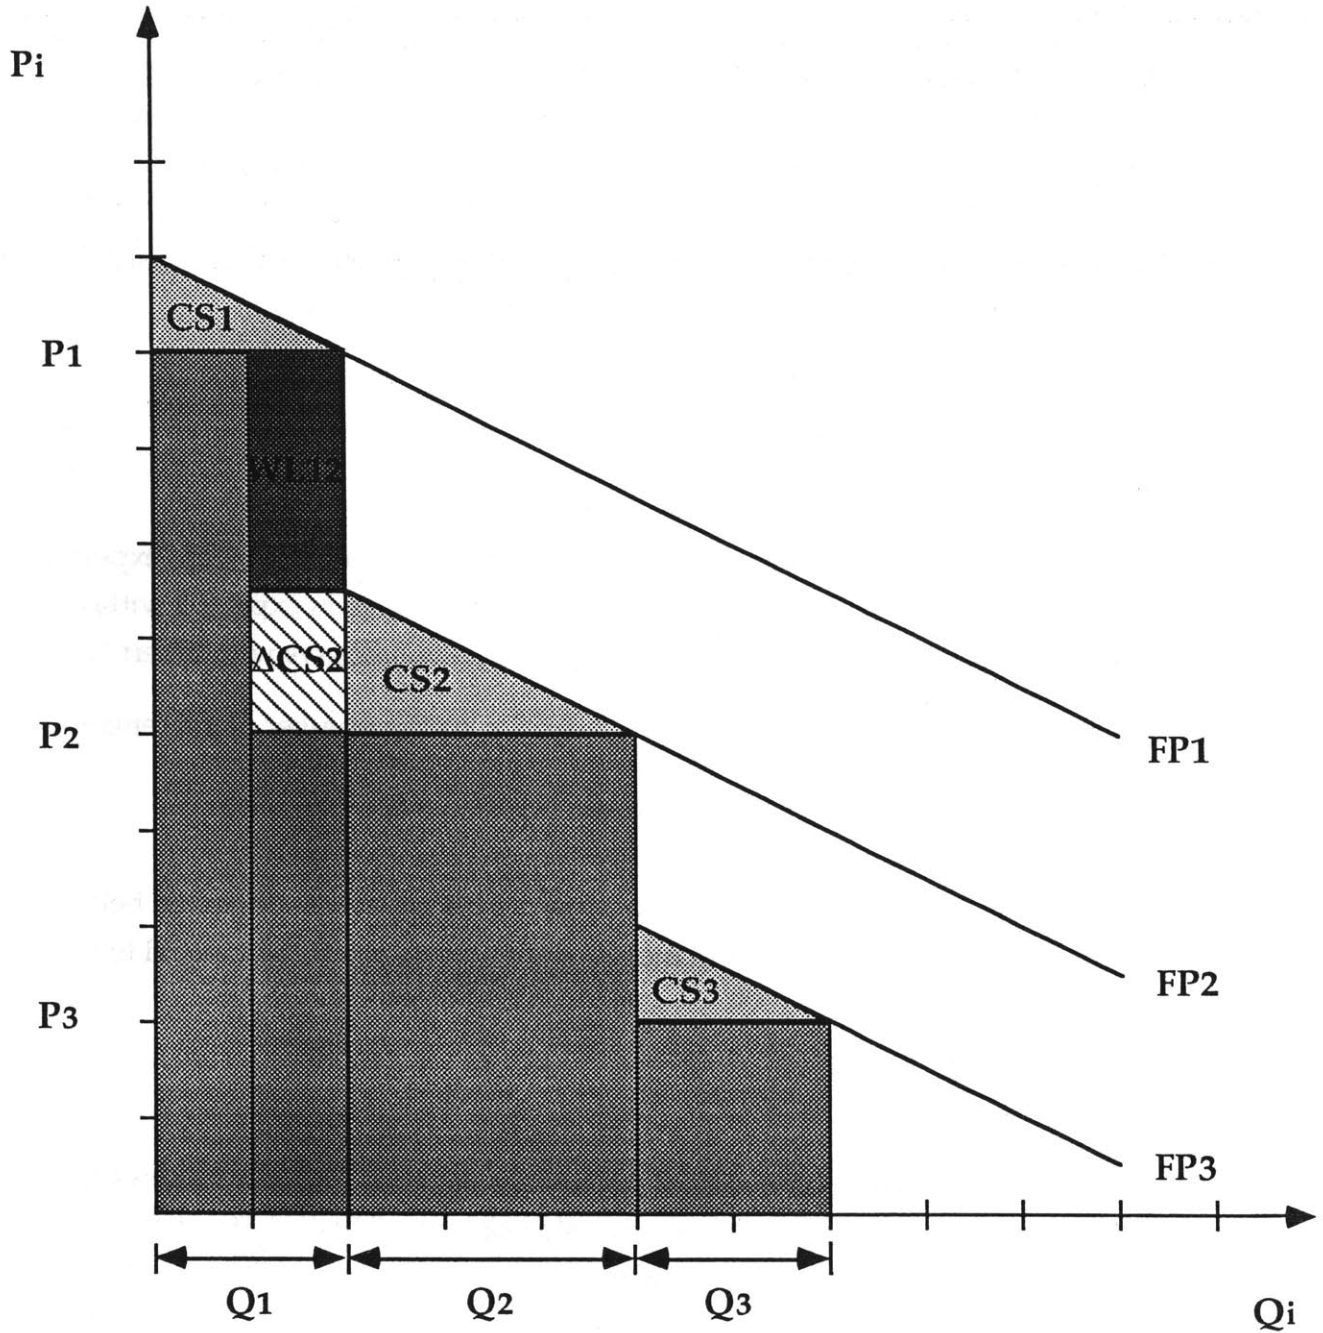# Consumer Surplus and Welfare Changes

Figure 4.9 - Generalized Cost Model With Diversion - Three Fare Product
Example - Consumer Surplus and Welfare Changes

There is an increase in the deadweight welfare loss experienced by the market when passengers divert. The welfare loss equal to :

$$\text{Welfare Loss}_{ij} = \left[\sum_{i<k\leq j} c_k\right] d_{ij} Q_i \qquad (4.25)$$

is associated with passengers diverting from less-restricted fare product i to more-restricted fare product j. The dark shaded box in Figure 4.9 shows the welfare loss associated with passengers diverting from fare product 1 to fare product 2 in the three fare product example. Both the welfare loss and consumer surplus gain totals for the OD market can be calculated by summing their respective values over all less-restricted fare products i and all more-restricted fare products j.

Varying Diversion Percentage by Exogenous/Endogenous Factors

If the diversion function is a function of exogenous factors, the diversion percentage can be calculated in advance and optimized in the same way as the fixed diversion percentages model. For example, the diversion percentages could be calculated as a function of the degradation costs of the fare products. Incorporating endogenous factors, such as the price levels of the fare products, in the diversion function would require that the optimization be performed with the diversion function explicitly input into the objective function or the constraint set. The same optimization techniques used to solve the constant diversion percentage could be employed to solve the modified endogenous diversion percentage problem.

By allowing the percentage of diverting passengers to vary systematically with price level, cross price elasticity effects can be incorporated in the model. Diversion percentages that are dependent upon the price differentials between the fare products occurring in the market can reflect, for instance, the greater number of passengers expected to divert as a result of an increased price discrepancy between the fare products. The diversion function between the fare products must be input explicitly into the objective function of the revenue maximization or added to the constraint set. The example of making the

diversion percentages dependent upon the (endogenous) price differential between the fare products would be :

$$\text{Max R} = \sum_{i=1}^{N} P_i(1- \sum_{j=i+1}^{N} d_{ij}(\Delta P_{ij}))Q_i + \sum_{i=1}^{N-1} \sum_{j=i+1}^{N} P_j d_{ij}(\Delta P_{ij})Q_i$$

where :

ΔP_{ij} = differential between the price level of fare product i and each lower-priced fare product j

In this case, the diversion percentages are a function of the endogenous price differential between the fare products of interest. Such a formulation allows for a very important aspect of passenger diversion behavior, the fare product price differentials, to be considered explicitly.

Nature of Costs Facing Diverting Passengers

The relationship between degradation costs and passenger diversion should be clarified as it relates to the generalized cost model. While diverting passengers may avoid segmentation by the fare product purchase restrictions in the generalized cost model, all passengers incur a cost associated with accepting the additional purchase restrictions of the less desirable airline fare products. Simply stated, the ability to divert need not lessen the penalty associated with accepting the more-restricted fare product; the purchase restrictions only fail to prevent the passenger from purchasing the discounted fare product. The flexibility of the generalized cost model formulation allows the cost incorporated to be positive, as motivated above, or zero, if passengers can meet the restrictions imposed by airlines at no cost.

The fare product degradation cost incurred by diverting passengers differs depending upon the form of the cost function assumed in the model. For example, in the constant cost model, the cost associated with accepting the restriction is simply the identical constant cost incurred by all passengers. In the increasing cost model, on the other hand, the degradation cost calculation is less straightforward. The cost incurred by each passenger can be calculated directly from the degradation cost function since the position of the passenger on all fare

162

product demand functions is assumed known. Thus, as long as the passenger can be identified by his place on the demand function, his perceived degradation cost can be calculated directly from the cost function.

If identification of the individual passengers and their associated degradation costs proves to be computationally prohibitive, assumptions can be made concerning the willingness to pay of the group of diverting passengers to calculate the expected degradation costs. For example, if the diverting passengers are assumed to be those with the lowest values of willingness to pay, the degradation cost can be assumed to be equal to the average value of willingness to pay within the interval of those more price-sensitive passengers. In any case, a diverting passenger degradation cost approximation can be made to make the model operational.

4.6.4 Impacts of Passenger Diversion

The model of demand currently used in most airline yield management systems groups passengers who actually purchase the discount products as members of the discount populations even though they may be willing to purchase higher-priced fare products. The revenue potential of passengers able to beat the segmentation schemes of airlines who are willing to pay higher amounts for air travel, if forced, are removed from consideration in potential market revenue calculations. As a result, such models systematically underpredict the revenue potential of the market. Diverting passengers are implicitly considered within the modeling of the independent fare product demand populations, ignoring their actual willingness to pay for air travel. In other words, the higher value of willingness to pay of such passengers is removed (not considered) when they are falsely assumed to be members of the discount population.

The basic framework of the generalized cost model (without diversion) presented here groups passengers by their willingness to pay. In the extended model, passengers are allowed to divert to lower-priced fare products than the airline intends in spite of their willingness to pay. The diverting passengers are initially identified as members of the higher-priced product population rather than the discount population. Thus, the revenue potential of the diverting passengers is explicitly identified in the generalized cost model. Even though the passengers

163

may not end up purchasing the fare product for which the airline targets them, the estimate of market revenue potential identifies their willingness to pay and purchase that fare product in the presence of perfect fencing capabilities or binding booking limits. Hence, a more correct representation of market revenue potential results using the generalized cost model than previously available in the literature.

4.6.5 Limiting Cases of Passenger Diversion Formulations

The limiting cases of passenger diversion provide insight into the strategies that an airline should adopt in the presence of ineffective segmentation devices between fare products. Consider the demand for differentiated fare products with diversion occurring between adjoining fare products only. As the diversion percentage for a single fare product approaches 100%, the demand for that product is driven to zero. All of the demand for a fare product would be lost to the next-lowest-priced fare product. In terms of the revenue recovered by the airline and the prices paid by consumers, it is as though the diversion-ravaged fare product were not offered on the market altogether. In a three fare product example for the constant cost model, the total diversion of fare product 1 consumers to fare product 2 is represented in Figure 4.10. The total shaded area in the figure represents the revenue for the airline without diversion while the darkly shaded area represents what revenue remains for the airline when diversion occurs.

The diverting passengers would incur the additional costs associated with accepting the more-restricted fare product resulting in a deadweight welfare loss for the market. The top portion of the lightly shaded box in Figure 4.10 represents the deadweight welfare loss to the market in the three fare product constant cost example. The extension to the N fare product constant cost case is straightforward. It is interesting to note that the airline would be wise to offer the unpurchased fare product at the price level of the more-restricted, neighboring fare product. This is true since demand would be greater for the less-restricted fare product as a result of the decreased degradation costs associated with it (provided that the segmentation mechanism was completely ineffective). The airline has added incentive to adopt this strategy since the

Generalized Cost Model With Diversion at 100% from Fare Product 1 to Fare Product 2
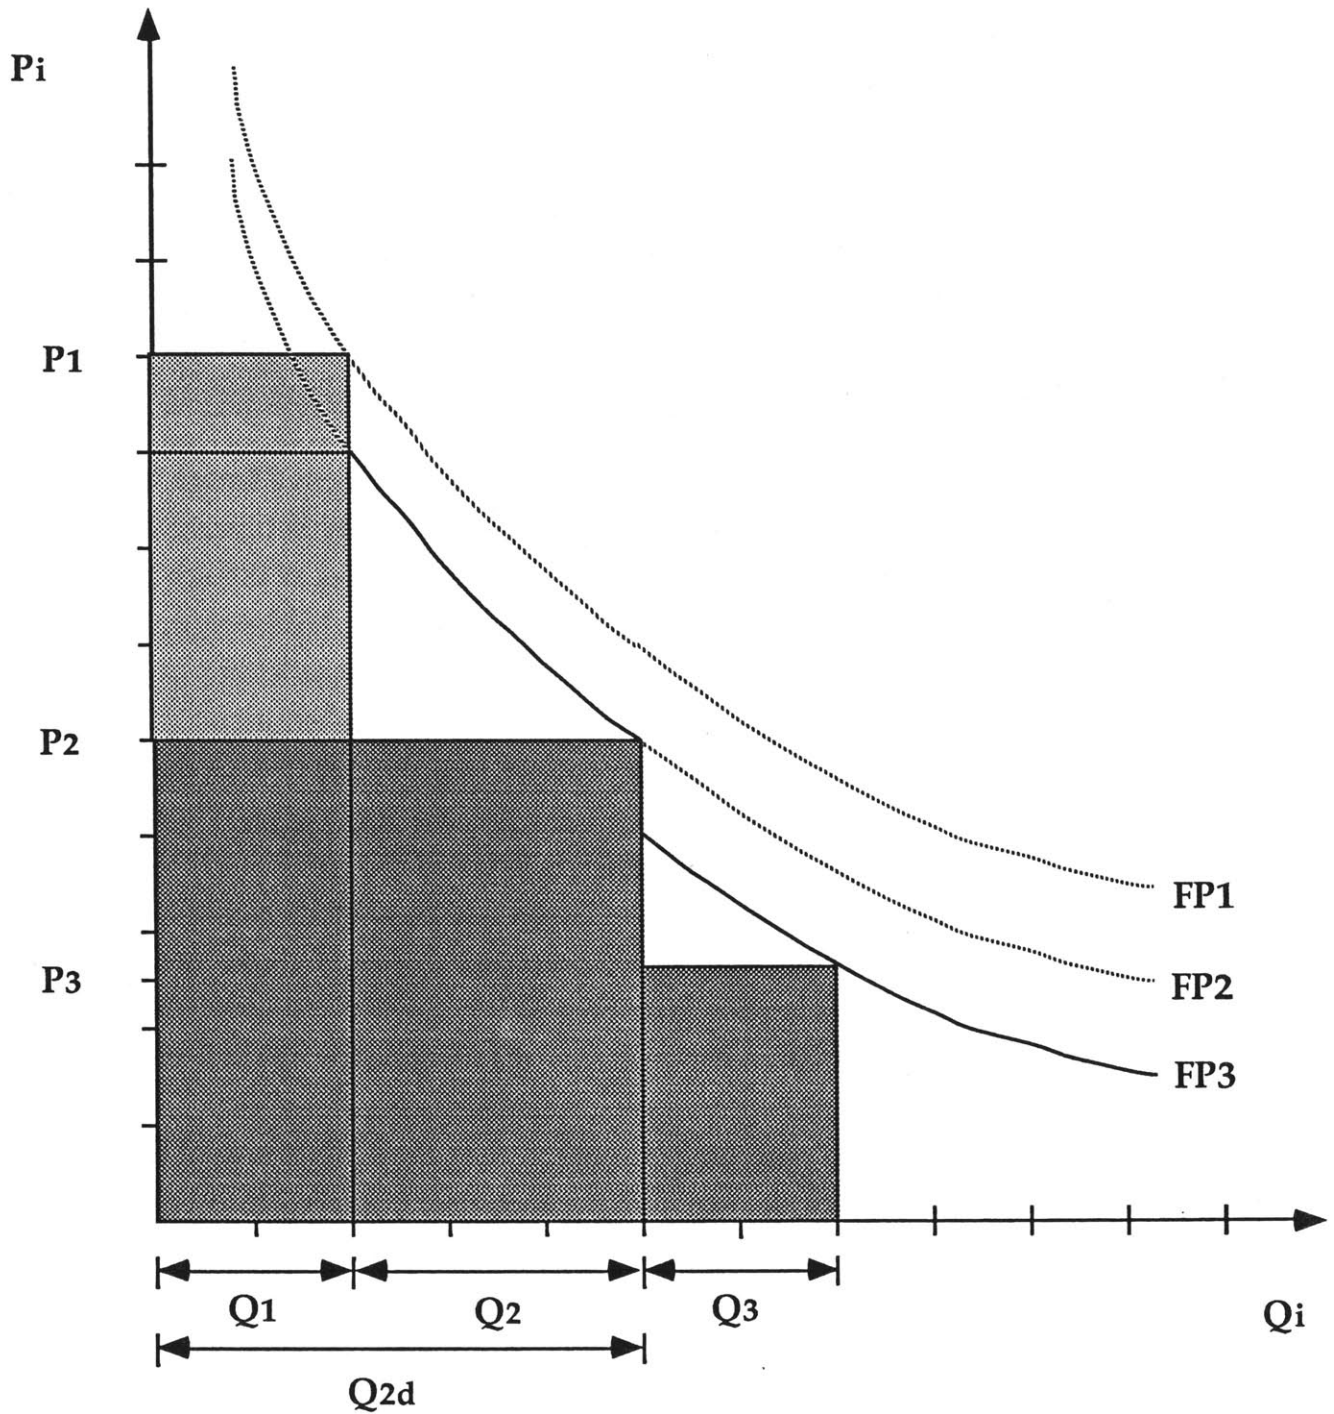

Figure 4.10 - Generalized Cost Model With Diversion at 100% from Fare
Product 1 to Fare Product 2

deadweight welfare loss does not affect consumer surplus, but rather, only airline revenue.

The deadweight welfare loss associated with the more-restricted fare product restriction bundle would be eliminated in the two fare product populations as well. The hatched area shown in Figure 4.11 represents the welfare loss incurred by the market as a result of offering both fare products 1 and 2 in lieu of simply offering fare product 1 when the fence separating the two fare products is completely ineffective. In the model forms other than constant cost, the lost welfare differs in amount but nonetheless is taken from airline revenue rather than consumer surplus.

The next limiting case of interest concerns completely ineffective fences for all fare products. In other words, all passengers have the ability to divert to the most restricted fare product. In this case, the market behaves as though there were a single fare product available on the market. Society, however, must bear the loss of welfare resulting from the perceived cost of the restrictions as a result of the single product being the least desirable. No such loss would occur if the airline offered a single unrestricted fare product, or if the cost to each passenger of accepting the more-restricted fare product were uniformly zero. Thus, in the presence of completely ineffective fences, it is to the benefit of the airline, consumers, and society to offer the unrestricted fare product in isolation instead of any restricted one. For the three fare class case, Figure 4.12 shows the increased number of passengers who choose to travel when the unrestricted fare product 1 is offered in isolation compared to (the most restricted) fare product 3.

With passenger diversion percentages equal to zero, the airlines would have the ability to perfectly identify and segment the consumer populations by their willingness to pay. The generalized cost model incorporating passenger diversion would mirror the model without diversion. The degradation costs associated with accepting the increased levels of purchase restrictions, however, would still affect demand. Clearly, this result is expected. It is interesting to note that from a total market welfare perspective, the no diversion case is the best possible result for N differentiated fare products considering degradation costs.

Generalized Cost Model With Diversion
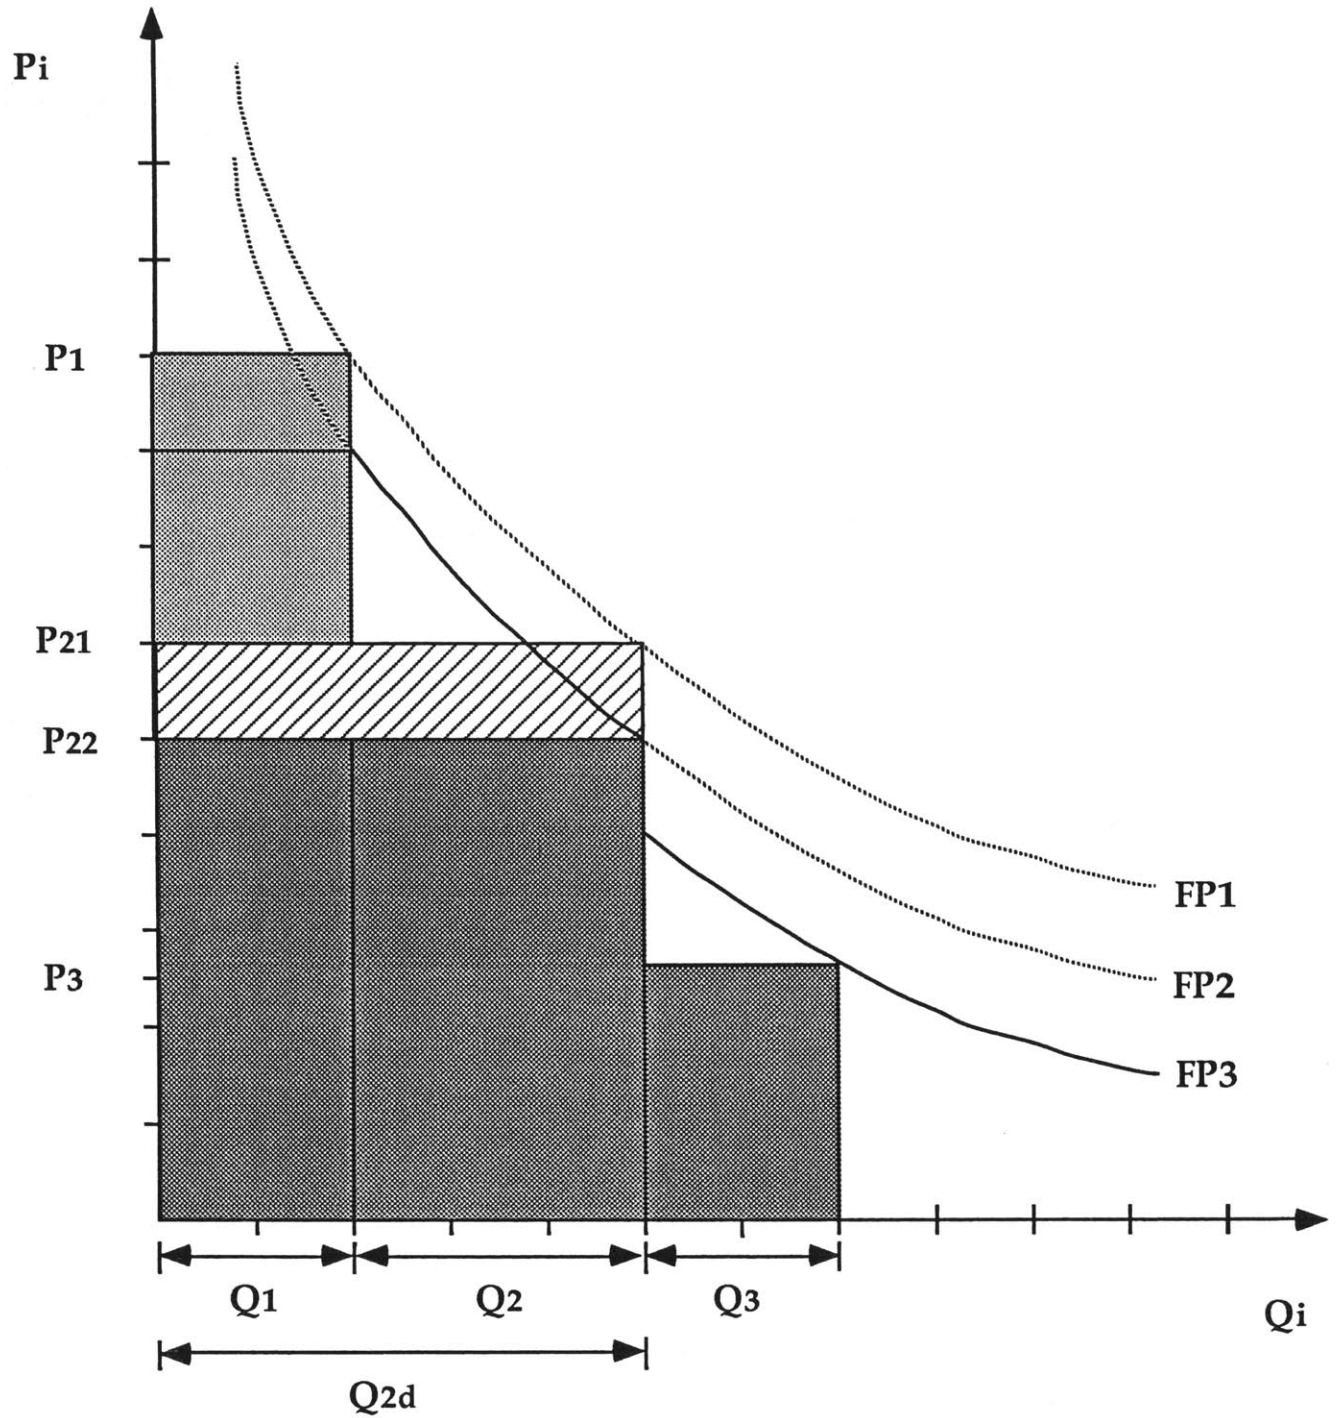# Welfare Loss with Ineffective Fences

Figure 4.11 - Generalized Cost Model With Diversion - Welfare Loss with
Ineffective Fences

Generalized Cost Model With Diversion
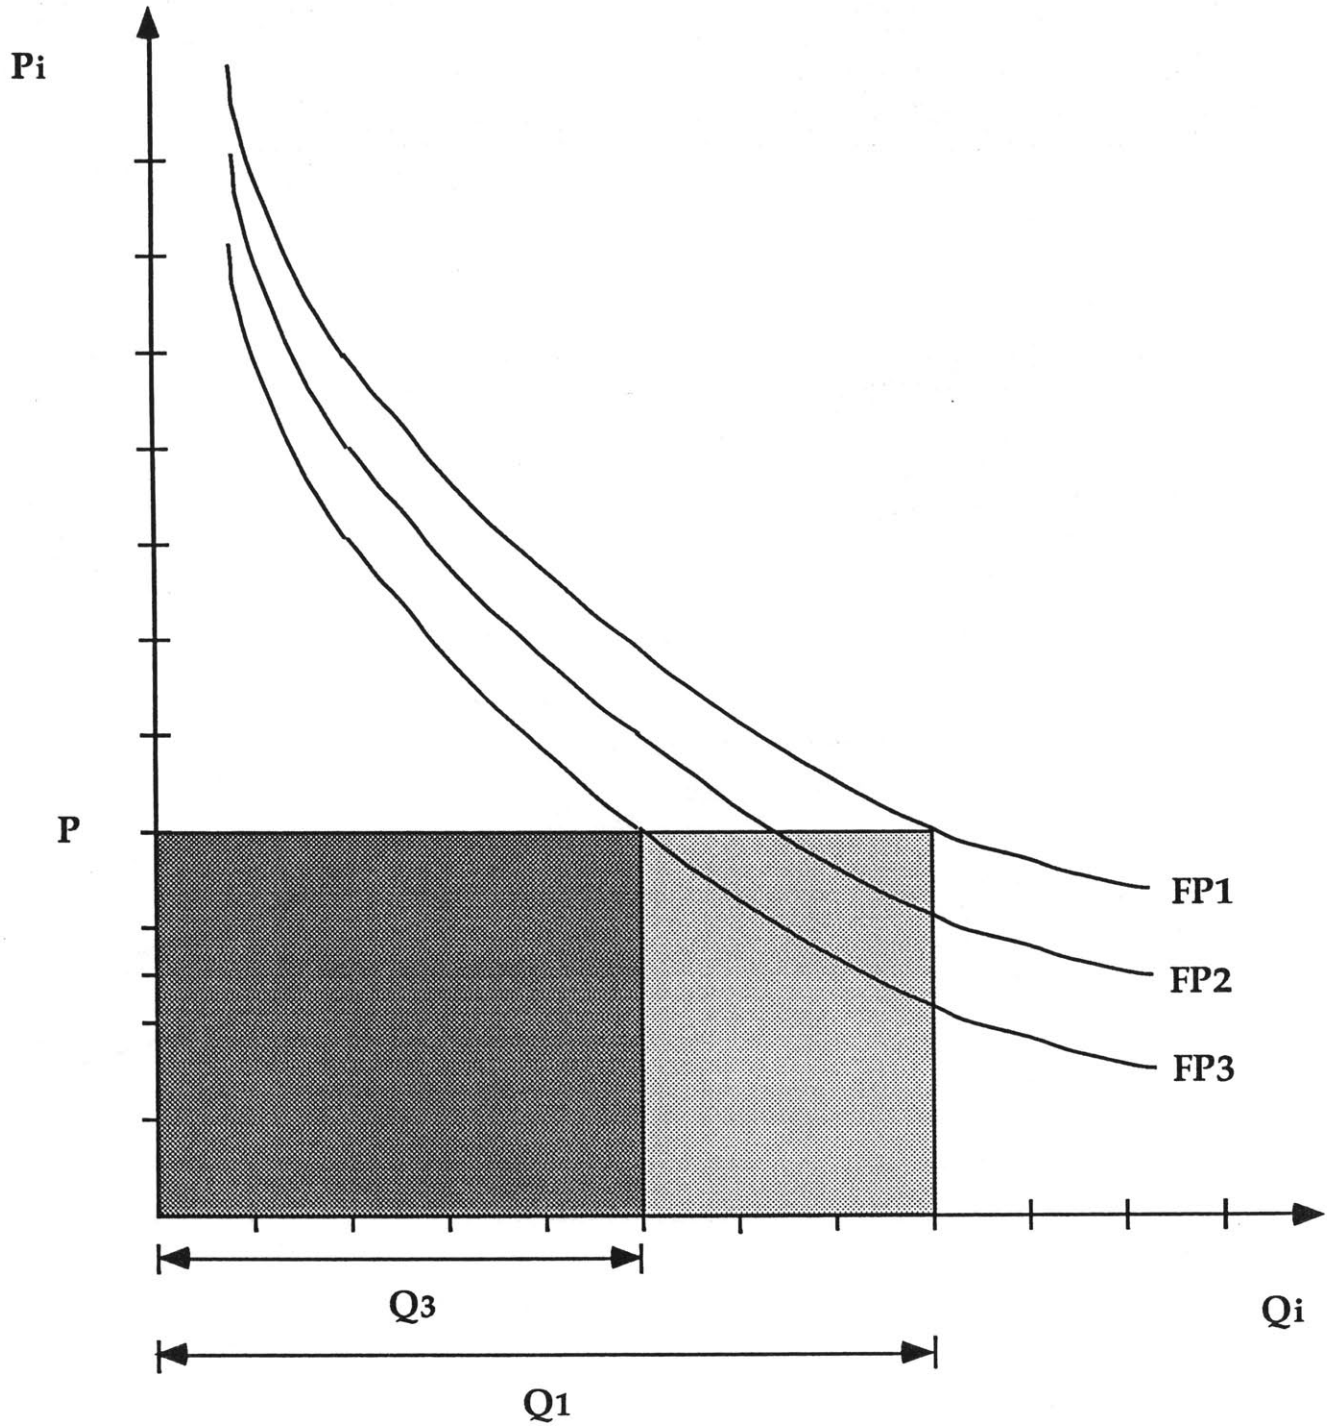# Single Fare Product Comparison

Figure 4.12 - Generalized Cost Model With Diversion - Single Fare Product
Comparison

Incorporating passenger diversion into the generalized cost model of airline fare product differentiation provides the first behaviorally motivated discussion of passenger switching behavior to appear in the airline yield management literature. Explicit consideration of passenger buy down behavior for the general N fare class case was performed for the first time. The incorporation of diversion represents a significant improvement over previous modeling attempts such as the standard monopoly price discrimination model presented in Chapter Two which relies heavily upon the assumption that airlines can perfectly identify and segment the consumer population by their willingness to pay. The relaxation of the perfect segmentation assumption more realistically models the situation facing airlines in the current industry environment. The application of booking limits to the fare products performed in the next section provides a similar behavioral motivation of passenger sell up not previously seen in the literature.

4.7 Application of Booking Limits to the Generalized Cost Model

The fixed capacity that exists for every aircraft directly affects the price levels that are charged by the carriers as well as the number of seats that are allocated to the different fare products. The capacity constraint can easily be incorporated into the generalized cost model by adding it to the constraint set. The optimal price levels can be determined through simple optimization techniques that reflect the effect of the capacity constraint on price levels. The capacity constraint is fixed for each departure based on the number of seats on the aircraft.

Applying booking limits to the fare products, however, is not so straightforward. The booking limits, in general, are set based upon the calculations of the yield management system as applied by the seat inventory control analysts of the airline. The booking limits differ by flight departure depending upon the demand forecasts and the other yield management system inputs. Each flight departure is subject to a different set of fare product booking limits making a simple inclusion in the optimization constraint set impossible.

The relationship between pricing and seat inventory control is one of the most interesting yet virtually unexplored areas in the revenue management field. The revenue maximizing seat allocations made by airlines are based upon the price

levels assumed for each fare class. Clearly, any model addressing airline fare product differentiation need consider the effects of booking limits and other capacity control devices used by airlines to correctly represent the realities of the current airline industry environment. Ideally, a joint price level/seat allocation optimization could be performed. The difficulties associated with introducing price level into the demand formulations addressed in the previous chapter, however, have been shown to be prohibitive. The introduction of booking limits to the generalized cost model provides useful insight into the interactions of price level and availability for airlines.

4.7.1 Booking Limits Motivation

Binding booking limits force passengers to reconsider their travel options because they make certain fare products and travel itineraries unavailable. A passenger who cannot purchase his preferred fare product/travel itinerary combination must choose an alternative action. The passenger may change fare products, travel itineraries, or both. In the extreme, the passenger may decide not to travel altogether. The case considered in the framework of the generalized cost model is a single flight departure in the OD market. Thus, the only options available to a consumer whose initial service request was denied are either to purchase a higher-priced fare product or decide not to travel.

Another fundamental assumption of the static generalized cost modeling approach is that all passengers arrive during a single booking period. This assumption presents no problems when supply and demand are equilibrated for each fare product. When capacity constraints are applied to the available fare products, however, the possibility that certain passengers will be denied their preferred fare products becomes quite real. Binding capacity constraints require that a method be chosen to allocate the limited resources among the passengers arriving in the single booking period. Although the static modeling assumption precludes an exact breakdown of passenger arrivals, bounding the maximum and minimum values attainable from the model can be done by making assumptions about the order of passenger arrivals within the single booking period.

To measure the effects of booking limits on the static model, an arrival process must be assumed within the booking period. For instance, passengers can be assumed to book in increasing or decreasing order of willingness to pay. Arrivals occurring in increasing order of willingness to pay would result in a lower bound on the number of passengers choosing not to travel when booking limits were applied. Conversely, arrivals in decreasing order of willingness to pay would provide an upper bound on passenger loss since the diverting passengers would arrive prior to those within the target population. A random arrival process within the booking period is also possible. In this case, an approximation to an average value of expected traffic loss from initial passenger service denials can be obtained.

Measuring the effects of the application of booking limits to the model also requires the establishment of rules concerning passenger behavior. The behavior of passengers who are denied a seat in their preferred booking class must be specified. The behavior of each passenger will be decided by his position on the isolated fare product demand functions and a simple assumption concerning sell up behavior. Clearly, a diverting passenger who is denied a lower-priced fare product is still willing to travel at the higher price (as demonstrated by his willingness to pay). Diverting passengers denied the lower-priced fare product are willing to sell up as high as the fare product to which they are initially targeted. Since diverting passengers are able to meet the restrictions of the fare product to which they divert, passengers initially denied a fare product are assumed to prefer the lowest-priced fare product that is available. Passengers who are not diverting, on the other hand, will choose not to travel in this framework since they are not willing to purchase any more expensive fare product than the one to which they have been targeted by the airline.

4.7.2 Booking Limits Applied to the Model

Under an assumed booking period arrival process, booking limits can be applied to estimate the impacts of yield management techniques on demand, revenue, and welfare under the behavioral assumptions of the generalized cost model. The booking limits can be applied to the fare products either exogenously after the price levels have been set or endogenously as a constraint in the joint price level optimization problem. The effects of booking limits on the market can then

be tested for both methods. Incorporating diversion in the model allows for the measurement of the effects of passenger sell up occurring after a fare product has sold out. Applying booking limits to the distribution of passenger demands allows us to calculate the effects of booking limits on passenger diversion. First, an example of an exogenous application of booking limits is illustrated.

Consider the three fare class example with booking limits applied to the middle fare product and diversion occurring between fare products 1 and 2. With prevailing market price levels P_1, P_2, and P_3, the darkly shaded area appearing in Figure 4.13 represents the revenue received by the airline. The lightly shaded area represents the revenue loss resulting from passenger diversion. When booking limits of $BL_2 = Q_2$ are applied to the flight, $(Q_{2d} - Q_2)$ passengers are denied their preferred fare product. Assuming that passengers arrive in decreasing order of willingness to pay, the $(Q_{2d} - Q_2)$ passengers denied fare product 2 will choose not to travel. The hatched area shown in Figure 4.14 represents the lost revenue associated with denying service to those passengers as a result of applying the booking limit BL_2. The revenue lost by the airline and the welfare lost by society as a result of passenger diversion (shown in the lightly shaded area of the figure) remain lost, even after the booking limits have been applied. Moreover, the total welfare associated with the passengers choosing not to travel (revenue and consumer surplus) is lost to the market when they are denied their preferred fare product.

Figure 4.15 shows the effect on revenue assuming that passengers arrive in increasing order of willingness to pay. In this case, the $(Q_{2d} - Q_2)$ passengers denied their preferred fare product (2) are those who wanted to divert to fare product 2 from fare product 1. Having been denied fare product 2, these passengers are compelled to purchase fare product 1 as initially intended by the airline. The airline recaptures the revenue represented in the lightly shaded area of Figure 4.15 as a result of the booking limits. The revenue increase, however, occurs at the expense of consumer surplus. In other words, the revenue recovery from the sell up behavior induced by the booking limits is equal to the consumer surplus that had been gained through the initial passenger diversion. The hatched area shown in Figure 4.15 represents the airline revenue shifted back from fare product 2 to fare product 1. In addition, the deadweight welfare loss resulting passenger diversion between the two fare products (shown in the upper
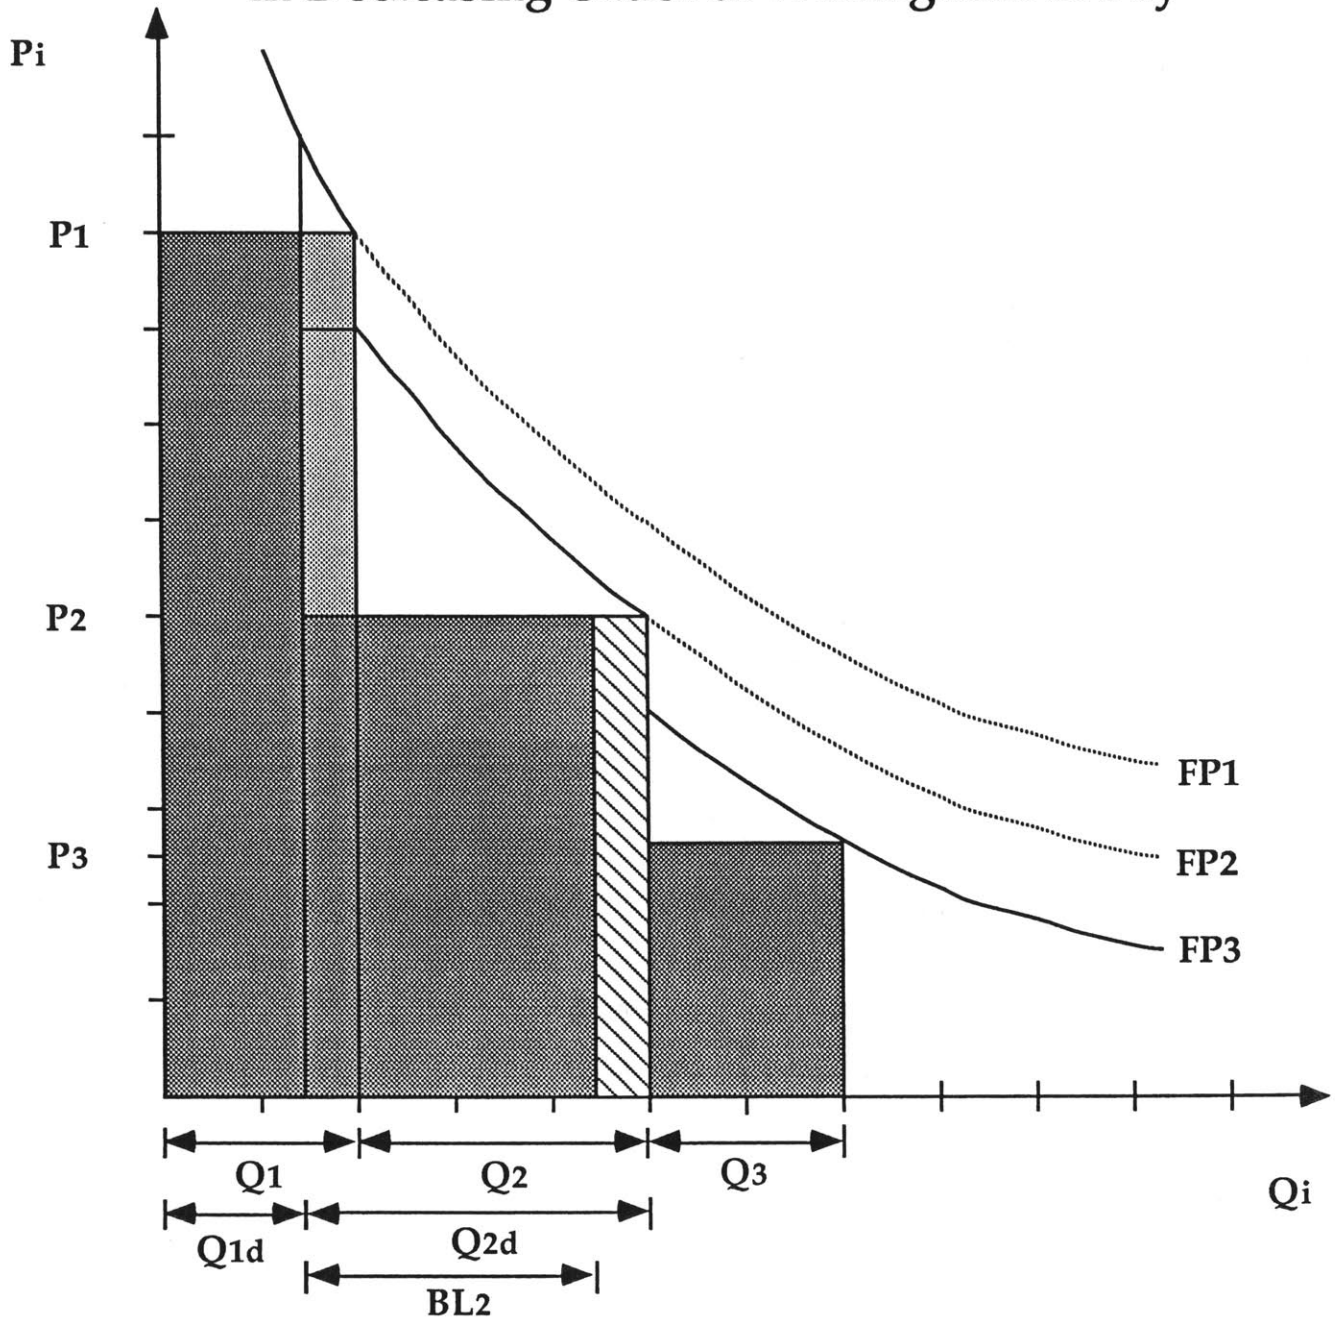
Figure 4.13 - Generalized Cost Model With Diversion - Three Fare Product
Example

Generalized Cost Model With Diversion
Three Fare Product Example
Booking Limits Application With Bookings
in Decreasing Order of Willingness to Pay

Figure 4.14 - Generalized Cost Model With Diversion - Three Fare Product
Example : Booking Limits Application With Bookings in Decreasing Order of
Willingness to Pay

174

Figure 4.15 - Generalized Cost Model With Diversion - Three Fare Product Example : Booking Limits Application With Bookings in Increasing Order of Willingness to Pay
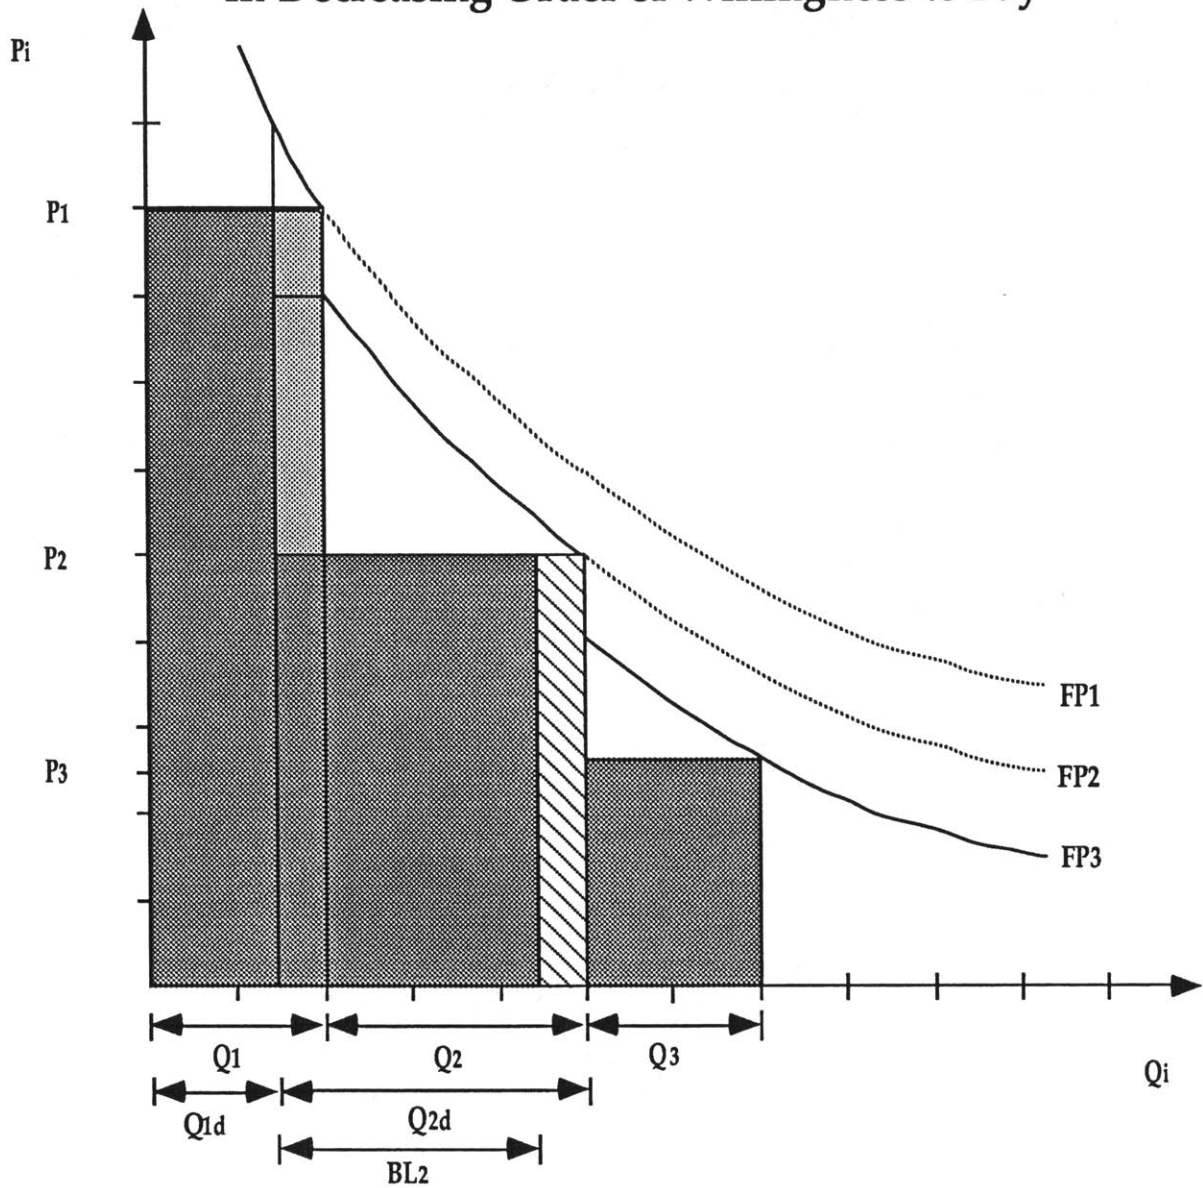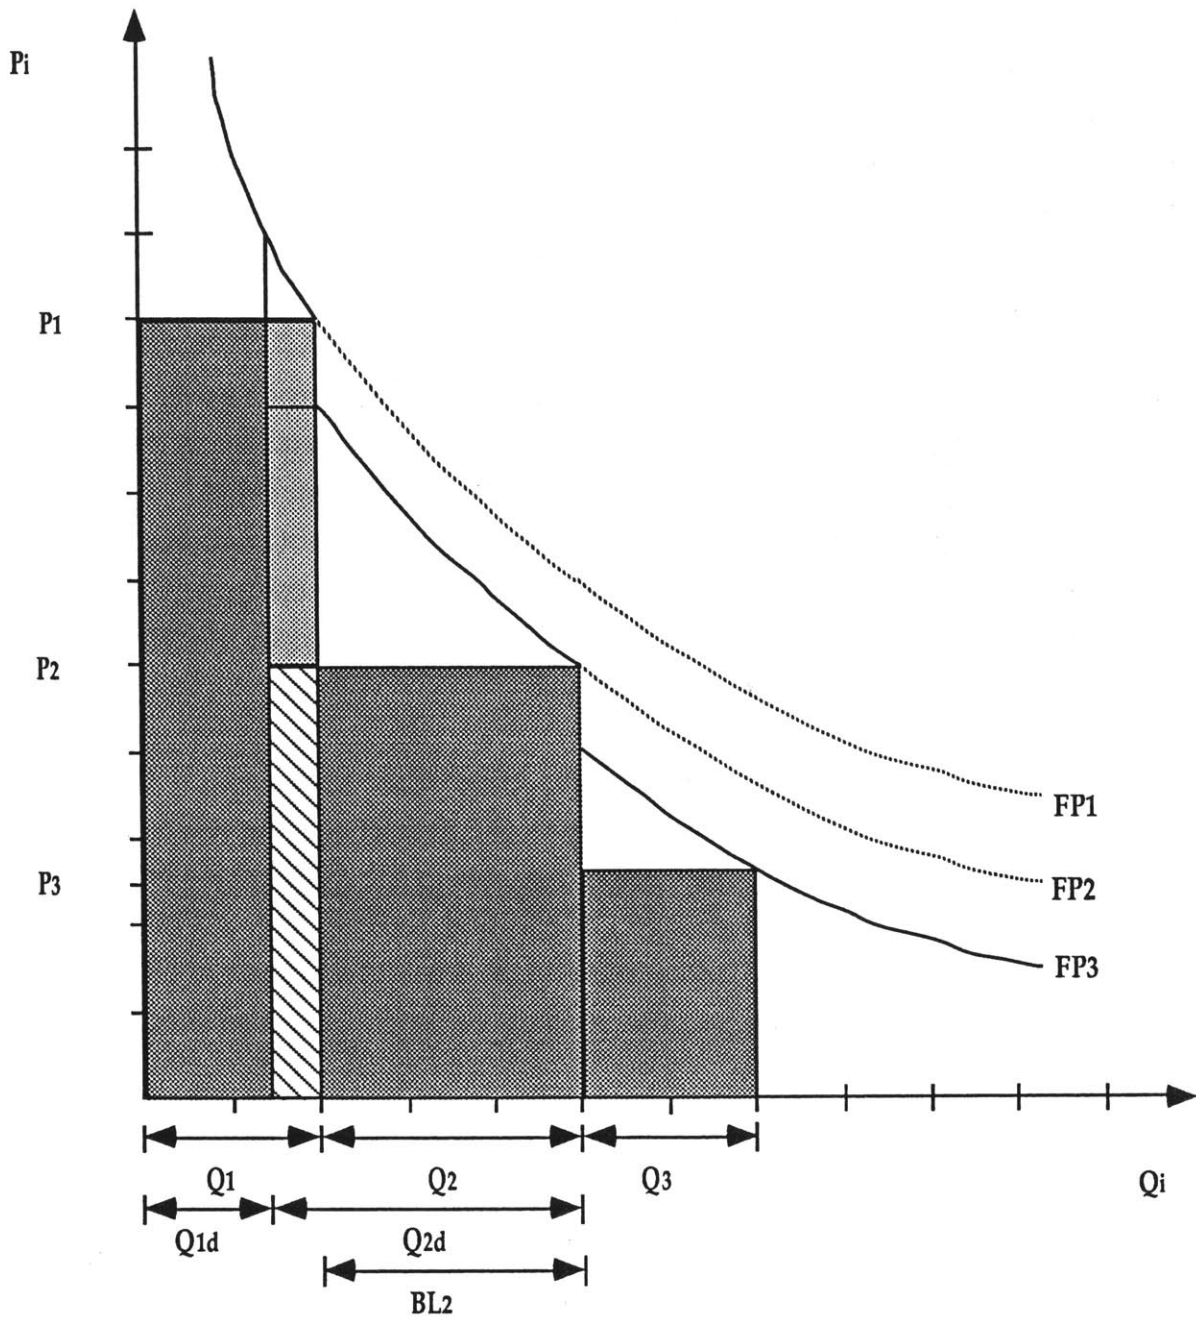
175

box of the lightly shaded area) is returned to the airline (and society) when the passenger sell up behavior is induced.

Clearly, the assumed arrivals in increasing and decreasing order of willingness to pay represent the upper and lower bounds, respectively, on market revenue change for a given set of booking limits. The example also demonstrates the indeterminance of the effect of booking limits on market revenue in the presence of diversion. The airline benefits from applying booking limit BL_2 when passengers with lower values of willingness to pay book first. Conversely, booking limit BL_2 is to the detriment of the airline when the opposite booking discipline prevails.

Endogenous applications of booking limits can also be performed using the generalized cost model. Adding booking limits to the constraint set in the joint price level optimization problem allows the testing of different cabin configurations under the prevailing demand conditions. Setting fixed capacity constraints for the first class, business class, and coach class cabins, for instance, allows the testing of the effects of various cabin configurations on optimal price levels and the resulting demand and revenue.

Service improvements can be captured in the generalized cost model in the degradation cost terms. Fixing the capacity of the aircraft and the cabin configurations, the effects of service improvements in the three cabins on the optimal price levels can be determined. In addition, the capacity and cabin configurations can be changed under a fixed fare product structure to evaluate different aircraft sizes and configurations on the optimal price levels and market revenues.

4.7.3 Booking Limits Summary

The behavior of the consumer population is more accurately modeled in the framework of the generalized cost model than the previous attempts in the literature because of the explicit modeling of passenger diversion. Prior literature has not rigorously motivated the reason that passengers are willing to upgrade to higher-priced fare products after being denied a discount fare product. The discussion of diversion presented identifies the passenger sell up

176

potential as simply the total number of diverting passengers. Direct calculation of the effects of passenger sell up behavior can be performed once booking limits have been applied to the model.

Because booking limits are the only large scale method to induce sell up, they can be viewed as damage control devices for airlines aimed at preventing the revenue dilution effects of passenger diversion that result from ineffective price discrimination techniques. An important purpose of booking limits is to induce passengers to pay an amount closer to their actual willingness to pay for air travel. The airline must be careful, however, not to discourage discretionary traffic stimulated exclusively through offering lower-priced fare products. In the absence of sell up considerations, booking limits only serve to maximize the revenue available from the mix of fare product populations on the aircraft through the use of yield management seat allocation optimization methods.

4.8 Chapter Summary

The chapter began by motivating the nature of airline demand in the context of the differentiated fare products currently offered on the market. The exact nature of the demand facing airlines was then built on the foundations of consumer perceptions about the cost of the restrictions placed on the fare products. The generalized cost model provided an estimable model of airline fare product differentiation which considers the decreased value of the more-restricted products explicitly when a functional form was assumed for the degradation costs. It also accounted for the interrelationships between the fare product price levels. The ability of the airline to perfectly identify and segment passengers by their willingness to pay was then relaxed to allow for passenger diversion. Finally, booking limits were applied to enable a more realistic evaluation of the model.

Chapter Five

Model Investigations Using the Generalized Cost Model of Airline Fare Product Differentiation

The number and types of fare products offered, or fare product structure, in an OD market influences the level of total passenger demand as well as the amount of revenue earned by the airline. The knowledge of how to set price levels and modify purchase restrictions to maximize revenue within the existing market fare product structure would provide an airline with an advantage in the area of revenue management. The generalized cost model of airline fare product differentiation presented in Chapter Four can provide an airline with a planning tool that gives insight into the effects that different price levels and purchase restrictions have on passenger traffic, airline revenue, or economic welfare, enabling more informed pricing policy decisions. In this chapter, various tests of the generalized cost model are performed which illustrate the relationships between fare product characteristics and market performance measures.

Airlines attempt to segment the consumer population into different market groups according to their willingness to pay. The multiple fare product structure now employed by airlines has resulted from second degree price discrimination or "self-selection" techniques used to maximize revenue (or best cover operating costs) through such a passenger market segmentation. The benefits and fallibilities of the self-selection techniques employed, however, are not well understood and have not been formally addressed in the literature. The fundamental relationships between fare product attributes, passenger demand, and airline revenue under the existing fare product structure are revealed using the generalized cost modeling framework which explicitly considers the market segmentation strategies employed by airlines.

The incorporation of passenger buy down or "diversion" into the generalized cost model provides a more realistic modeling of the market demand behavior actually facing airlines. An understanding of the effects that passengers who avoid airline market segmentation schemes have on demand and revenue can help the airline to formulate pricing strategies (or apply capacity controls)

designed to counteract the revenue dilution of passenger diversion. Furthermore, variations in the passenger diversion parameters input to the generalized cost model allows the airline to test the sensitivity of the analysis results to the level of revenue dilution.

Airlines have the ability to vary certain factors of supply as well as fare product characteristics to improve revenue performance. For instance, the purchase restrictions accompanying the fare products offered in the market may be changed to increase revenues when market price levels are fixed. Alternative fare products can be compared, using the generalized cost model, to determine which one is expected to provide the greatest benefit to the airline in the prevailing OD market environment. In addition, the model can be used to quantify the effects of variations in the aircraft capacity and configuration serving a market, offering airline planners a method to increase the revenue earned in a particular market. An understanding (and quantification) of the effects that supply factors under the control of the airline have on optimal price levels and revenue allows marketing planners to better set price levels or modify supply to maximize revenue, making the generalized cost model a powerful airline planning tool.

The generalized cost model also quantifies several tradeoffs that airlines face, for example, demand stimulation versus revenue dilution. In the model, the increased (degradation) costs facing consumers associated with accepting more purchase restrictions can be traded off against a higher rate of passenger diversion occurring between the fare products to provide a comparison of the revenue impacts of each effect. The impacts of introducing fare products designed to stimulate demand can be analyzed with proper consideration given to the amount of revenue dilution expected to occur as a result. More informed decisions concerning fare product introductions designed to stimulate discretionary air travel demand can thus be made. Among the other tradeoffs facing airline planners that are tested in this chapter using the generalized cost model are the costs and benefits of proposed service improvements and the problem of aircraft cabin configuration.

The chapter begins by describing how the generalized cost model can be used to evaluate the demand and revenue expected in the market for a given set of price

179

levels. Endogenously determining price levels through a joint optimization of all fare product price levels using the constant cost formulation of the generalized cost model provides a measure of the maximum amount of revenue that can be earned in the current market environment. The fundamental relationships between optimal price levels and supply factors such as the severity of fare product purchase restrictions or the configuration of the aircraft are then discussed based upon the constant cost model outputs using both analytical and numerical techniques. Finally, different airline pricing policies are tested to demonstrate the uses of the model as a decision support and planning tool for airline management.

5.1 Analytical and Numerical Analysis Using the Generalized Cost Model

The primary application of the generalized cost model is to analyze and provide suggestions for the design and modification of the fare product structure in a specific air transportation OD market. In a deregulated environment, airlines exercise control over the price levels and purchase restrictions in all OD markets to the extent that competition from other carriers does not prevent such control. To this end, airlines seek to offer the combination of prices and restrictions that best enhance their revenues. Consequently, the question of how to best set price levels is of great interest to airlines.

The impacts of pricing decisions vary depending upon the existing OD market conditions. Market conditions include the passenger price level-demand relationship, perceived fare product degradation costs, available seat supply, and the ability of the airline to segment the consumer population. Comparative statics analysis techniques using the generalized cost model can measure the effects of changes in the fare product structure or other market conditions on demand, revenue, and welfare. In particular, the analysis performed in this chapter seeks to answer the following questions :

1) How can prices be set under fixed market conditions to improve results?
2) How do market condition changes affect optimal price levels?
3) How can OD market conditions be changed to improve results?

All analyses are performed from the viewpoint of the airline. Commercial airlines are interested in making the greatest possible return on investment and, thus, seek maximum profits. With a fixed aircraft allocation and in turn fixed operating costs, this amounts to maximizing revenues. Accordingly, the objective functions of the mathematical optimization formulations presented are limited to airline revenue maximization with price levels as decision variables. Other objective functions may be substituted to address different viewpoints (*e.g.* that of a regulatory body); however, such analyses are left to future research efforts.

The power of the model as a planning tool is revealed when it is used to measure the effects of airline pricing and fare product attribute decisions on traffic, revenue, and welfare in the OD market. Different airline pricing strategies can be evaluated based upon their expected impacts on revenue and other market conditions. Comparisons of model results under different expected market conditions test the sensitivity of different pricing and fare product design policies to changes in the OD market environment. In this way, the relationships between degradation costs, price levels, demand, revenue, and welfare under different airline OD market conditions can be identified.

Price levels can either be set exogenously or determined through the revenue maximizations and used to evaluate pricing strategies or initiatives. The impacts of changes in the OD market fare product structure can be quantified through a simple evaluation of existing market conditions using the generalized cost model with exogenously determined price levels. A joint optimization of price levels for all fare products provides insight for airlines concerning pricing decisions and their effect on market revenue potential.

The flexibility of the generalized cost model provides the ability to test numerous behavioral assumptions about demand. This flexibility emanates from the many versions of the model that can be tested. Different functional forms of demand thought to characterize markets can be evaluated. Assumptions about the nature of the fare product degradation costs can also be tested. In addition, the passenger segmentation assumption can be varied to offer a more realistic view of passenger behavior. The result is a planning tool with the ability to test the

relationships most important to pricing decisions in an OD market under varying scenarios.

In summary, the effects of fare product attributes (*e.g.* price level) and market condition changes under a stable fare product structure can be quantified using the generalized cost model. The conditions existing in the OD market under study must first be characterized, however, before any analysis can be performed. The results of the model, under the assumed behavior of passengers considering travel in the market, can then be used to quantify the impact of the pricing strategies that an airline seeks to employ. Airline planning decisions can be improved with the results obtained by using the generalized cost model to develop a better understanding of the underlying relationships between price levels, restrictions, demand, and revenue.

5.1.1 Inputs to the Model

The generalized cost model provides a flexible characterization of air travel demand under different price levels and fare product structures. The model requires the characterization of several factors influencing travel behavior in the OD market to be used as inputs. The inputs are required to characterize passenger behavior so that the results of the generalized cost model can be evaluated under plausible assumptions. The major inputs to the model are :

1) Demand functions for each fare product
2) Degradation cost functions for each additional restriction bundle
3) Passenger segmentation assumption
4) Price levels for each fare product
5) Aircraft capacity for the OD market

As discussed in the model presentation of Chapter Four, a required input to the optimization model is the functional form of the perceived degradation costs facing consumers due to fare product purchase restrictions. The form of the costs depends upon the assumed sensitivities of the travel population. Within the context of the joint price level optimization, the perceptions of consumers considering air travel are of interest to the airlines in the quest to maximize the revenues received from providing air travel services.

182

The selection of a specific functional form of degradation costs facing consumers determines which form of the generalized cost model should be used for analysis. For instance, selection of a constant degradation cost facing consumers leads to the constant cost model being employed. The model allows for computational testing using any intuitively pleasing functional form of degradation costs. A practical model implementation, however, may limit the functional form to an estimable one. The assumption and specification of a functional form of degradation costs is equivalent to the selection of a model form and calibration, respectively.

The price-demand relationship existing in the market for the potential air travel consumers must be defined to properly measure the effects of the fare product structure on the OD market. In particular, a functional form for demand must be assumed in order to yield an operational version of the generalized cost model. The functional form of demand chosen should accurately reflect existing market conditions. The price elasticity of demand facing different consumers, for instance, can determine the appropriateness of a given demand formulation. Constant elasticity demand functions can be selected if passengers are thought to have a uniform sensitivity to price level changes. Linear demand functions, on the other hand, imply a measure of price elasticity that changes with willingness to pay. Both linear and constant elasticity demand functions are used as inputs to the generalized cost model in this chapter.

Next, assumptions about the ability of the airline to segment consumers into distinct fare product populations must be made. Most existing yield management seat allocation models assume the ability of the airline to perfectly segment the demand population into distinct and independent fare class (product) populations. In the context of the generalized cost model, a perfect segmentation assumption is only one alternative. A valuable extension of the model allows a functional form of passenger buy down, or diversion, behavior to be chosen. Thus, the ability of passengers to avoid airline segmentation schemes can be incorporated explicitly. Moreover, variation of the segmentation assumption allows for the evaluation of different market condition scenarios when passenger diversion occurs.

The aircraft capacity serving the OD market must also be identified. The existence of a well defined capacity constraint depends upon the scope of the model being considered. If the generalized cost model is being used to analyze an OD market with multiple flights, the representation of the capacity constraint is ambiguous. If, however, the OD market is served by only one flight or the model is being used to evaluate the demand for a single flight, then a capacity constraint can be applied unambiguously. The examples presented in this chapter are limited to the single flight/OD market level making the capacity constraint equivalent to the number of seats on the aircraft serving the market.

An unambiguous capacity constraint also allows for booking limits on each fare product to be applied to the solution. In other words, meaningful booking limits may only be applied to a model whose scope of analysis is limited to the single flight/OD market case. The analysis of booking limits shows the impacts that capacity constraints on the individual fare products have on traffic and revenue. The application of booking limits to the solutions allows for the measurement of passenger sell up in the market when passenger segmentation is imperfect.

If a single flight serves the market, then both booking limits and a capacity constraint can be applied to the analysis. The evaluation of the effects of the capacity constraint or booking limits requires the assumption of a booking arrival discipline within the single booking period (assumed in the generalized cost model). Without such an assumption, the effects of limits on capacity are likely to be indeterminate. As previously stated, passenger booking inquiries arriving in increasing, decreasing, or random order by willingness to pay are the candidate processes. Evaluation of the model under the different arrival scenarios provides a more precise look at the effects of price levels and restrictions on the OD market under availability constraints.

The price levels must be specified to complete the market characterization and enable an analysis of passenger demand, revenue, and welfare. Price levels can either be determined exogenously or endogenously and input to the model. The exogenous calculation of price levels may be the result, for example, of a competitive equilibrium or of price levels dictated by a regulatory agency. The endogenous setting of price levels, on the other hand, can be done using non-linear optimization techniques applied to the generalized cost model. In either

184

case, the price levels are used as input to the generalized cost model to determine the expected demand and revenue in the market under the prevailing conditions. A more complete discussion of price level determination in the generalized cost model appears next.

5.1.2 Exogenous Price Level Determination

In the case of externally determined price levels, the evaluation of the existing environment facing airlines requires the ability to perform analysis given the current market conditions. Such a situation may result, for instance, when a regulatory agency dictates price levels for the OD market. The generalized cost model provides the airline with the ability to evaluate the expected demand, revenue, and welfare for the market under the prevailing market price levels. Therefore, the current state of the market can be quantified and used as a base measure for planning purposes. The demand and revenue figures that are expected from the prevailing market condition analysis can be compared to the optimal values obtained using optimization techniques to measure the percent difference from optimal of existing conditions.

For a general N fare product structure, there are N published fares P_i for all $i \in N$ existing in the market. These price levels will determine the number of passengers who purchase each of the N fare products under the assumed price level-demand relationship existing in the market. The generalized cost model attempts to quantify the level of demand for each fare product given the prevailing OD market conditions (including price level). The resulting revenue requires only a trivial calculation as highlighted in the Chapter Four examples.

5.1.3 Price Level Determination by Joint Price Level Optimization

While individual airlines are unable to dictate the price levels in the majority of their markets, each has the ability to influence the price levels through the publication of fare products in the computer reservations systems. Thus, it is of interest to the airline to understand the benefits and pitfalls of raising or lowering price levels. To gain insight into the effects of price levels on airline revenue and societal welfare, the joint price level optimization can be used. The formulation of the joint price level optimization under the assumption that the

185

airline can perfectly identify and segment the population by their willingness to pay appears below :

$$\text{Max } R = \sum_{i=1}^{N} P_i Q_i$$

Subject to :

$$\sum_{i=1}^{N} Q_i \leq \text{Cap}$$

$$Q_i \geq 0 \qquad \text{for } i = 1,...,N$$

where | R | = | total revenue |
|---|---|---|
| Q_i | = | seats allocated to fare product i |
| P_i | = | average fare charged to fare product i |
| N | = | total number of fare products |
| Cap | = | total aircraft capacity |
| c_k | = | cost associated with accepting the restriction bundle accompanying fare product k |

The optimization requires the joint calculation of price levels for all fare products to present the best suggestion for setting price levels in the N fare product OD market under an assumed specification of the price level/demand relationship. The optimal price levels serve as a performance benchmark to which the prevailing market fares (or any other set of fares) can be compared. The model suggests the direction in which price levels should be moved to improve market revenue performance. The results from the optimal price levels can be compared to those from the prevailing price levels to demonstrate the proximity of current conditions to optimality. The analysis also provides information about the proximity to optimality of any other set of price levels.

The joint price level optimization problem has similar requirements to a simple evaluation of the generalized cost model. Essentially, the generalized cost model must be fully specified (with the exception of price levels) to enable the price level optimization. The constant cost formulation is the form of degradation cost

186

function used to demonstrate the optimization here. Additionally, the optimization formulation requires a functional form of passenger demand to test the model. In this research, two general forms are used, linear and constant elasticity. The analysis must be limited to the single flight/single market case to preserve the integrity of the capacity constraint. Price levels are the decision variables and, thus, are determined as outputs of the joint optimization. The evaluation of passenger demand and revenue along with welfare calculations can be made once the price levels have been determined.

The analysis presented here is identical whether or not the price levels have been determined optimally or by any other method. The analysis shown focuses on the results of the joint price level optimization applied to the generalized cost model. The results of applying the prevailing (or any other) price levels to the model would produce similar results and is not shown except where insightful or when optimization is not possible. Nonetheless, the flexibility to evaluate the model under any prevailing set of price levels should be noted.

5.1.4 Outputs of the Model

The expected values of the following measures can be determined from a fully specified generalized cost model :

1) Demand
2) Revenue
3) Consumer surplus
4) System welfare

Preferred airline strategies can be obtained by testing proposed changes to the fare product structure using the generalized cost model under the conditions thought to represent the current OD market environment. The inputs and outputs of the generalized cost model can also be used to develop insight into the underlying relationships between the factors influencing airline travel demand. The traffic and revenue calculations are unambiguous for all versions of the generalized cost model. The consumer surplus and welfare calculations require behavioral assumptions about the passenger booking arrival process within the single booking period to be identified.

In summary, the generalized cost model requires several inputs including a functional form for demand, the degradation costs facing consumers, and a passenger segmentation scheme. The OD market must be characterized with the inputs which contain assumptions about the conditions existing in the market. The generalized cost model may then be used to evaluate the market conditions in terms of expected passengers, revenue, and welfare. Testing of the specific functional forms and practical applications of the generalized cost model comprises the balance of the chapter.

5.2 Constant Cost Model with Linear Demands

Chapter Four introduced forms of the generalized cost model that could be used to characterize the perceived degradation costs facing consumers considering air travel. In this chapter, the constant cost model is combined with assumed functional forms of demand for analytical and numerical analysis. The resulting model specifications are used to demonstrate the effects of changes to the fare products in the OD market under different behavioral scenarios. Although most results are illustrated using a three fare product example, all can be generalized to the N fare product case.

The constant cost model with linear demand functions will serve as the base case model formulation against which the majority of comparisons are made. Although any of the model forms could be used for this purpose, the linear constant cost model has been chosen for its expository simplicity. To uncover the underlying relationships between the market characteristics and the optimal price levels, obtaining a general closed form analytical solution is preferable. Thus, where possible, such analytical solutions are provided. The simple formulations of the generalized cost model are amenable to analytical solutions for the N fare class case and appear first. Numerical examples are provided for the more complicated model extensions.

The Lagrange multiplier method was used to determine the optimal solution to the general N fare product joint price level optimization. The solution to the Lagrange multiplier method was checked to determine whether or not the necessary conditions of the constrained optimization problem were met. The

correct result could be easily identified from the sum of the number of passengers purchasing each fare product at the optimal price levels. That is, when the total number of passengers is equal to capacity, the Lagrange multiplier method with a binding capacity constraint provides the optimal solution. When the capacity constraint does not bind, the Lagrange multiplier term drops out and the result is a simple unconstrained optimization.

When analytical techniques prove too cumbersome, the joint price level optimization can be solved using non-linear optimization techniques. For instance, the optimal solutions to the program can be provided using off-the-shelf software packages (*e.g.* GAMS) which perform non-linear optimization. The more complicated forms of the generalized cost model are solved numerically using the GAMS package under plausible OD market characteristics.

5.2.1 Capacity Unconstrained Case

Unconstrained optimization yields the optimal price levels for the joint optimization when the capacity constraint is not binding. In other words, if the aircraft is not full, unconstrained optimization is the appropriate methodology. The objective function used (assumed to be from the perspective of the airline) is a simple revenue maximization. Other objectives could be tested, however, since airlines are ultimately responsible for the pricing of their services, revenue maximization seems appropriate.

The basic constant cost model (without diversion) has been solved analytically using the techniques of unconstrained optimization (or, equivalently, the Lagrange multiplier method with $\lambda = 0$) assuming a linear demand function. The demand function for the linear case appears as follows for the constant cost model :

$$P_i = P_0 - a\sum_{j \leq i} Q_j - \sum_{j \leq i} c_j \qquad (5.1)$$

Substituting the assumed linear demand functions into the objective function of the joint price level optimization yields the following simplified objective :

189

$$\text{Max } R = P_0 \sum_{i=1}^{N} Q_i - a \sum_{i=1}^{N} Q_i^2 - a \sum_{\forall (i \neq j)} Q_i Q_j - \sum_{i=1}^{N} \sum_{j=1}^{i} c_j Q_i \qquad (5.2)$$

Using simple calculus, the unconstrained optimality condition for each fare product i is :

$$\frac{dR}{dQ_i} = P_0 - 2aQ_i - \sum_{j \neq i} Q_j - \sum_{j=1}^{i} c_j = 0 \qquad (5.3)$$

The unconstrained optimal solution in terms of the number of passengers requesting service using each fare product i is :

$$Q_i^* = \frac{\left[P_0 - \sum_{k=1}^{N} k c_k \right]}{(N+1)a} + \frac{\sum_{k=i+1}^{N} c_k}{a} \qquad (5.4)$$

A substitution of the optimal fare product demand levels into the demand functions yields the optimal price levels for each fare product :

$$P_i^* = \left[1 - \left(\frac{i}{N+1} \right) \right] \left(P_0 - \sum_{k=1}^{i} k c_k \right) + \left(\frac{i}{N+1} \right) \sum_{k=i+1}^{N} k c_k - \sum_{k=i+1}^{N} i c_k \qquad (5.5)$$

The solution to the unconstrained optimization, of course, is only valid in cases where $\sum_{i=1}^{N} Q_i \leq Cap$.

5.2.2 Capacity Constrained Case

If the aircraft capacity constraint is binding, the Lagrange multiplier technique provides an analytical solution to the non-linear joint price level optimization program with linear demand functions. The Lagrange multiplier method presented assumes that the aircraft capacity constraint is binding and incorporates that constraint into the objective function with a multiplier term. Substitution of the linear demand functions and capacity constraint into the

revenue maximization objective yields the Lagrangian for the capacity constrained optimization :

$$\text{Max } \mathcal{L} = P_0 \sum_{i=1}^{N} Q_i - a \sum_{i=1}^{N} Q_i^2 - a \sum_{\forall\,(i \neq j)} Q_i Q_j - \sum_{i=1}^{N} \sum_{j=1}^{i} c_j Q_i - \lambda \left[\text{Cap} - \sum_{i=1}^{N} Q_i \right] \qquad (5.6)$$

Differentiation of the Lagrangian with respect to each fare product demand level along with the capacity constraint provides the capacity constrained optimality conditions :

$$\frac{d\mathcal{L}}{dQ_i} = P_0 - 2aQ_i - \sum_{j \neq i} Q_j - \sum_{j=1}^{i} c_j - \lambda = 0 \qquad (5.7)$$

$$\frac{d\mathcal{L}}{d\lambda} = \text{Cap} - \sum_{i=1}^{N} Q_i = 0 \qquad (5.8)$$

Solution of the program through back substitutions of the optimality conditions yields the optimal demand levels for the capacity constrained case :

$$Q_i^* = \frac{\text{Cap}}{N} - \left(\frac{1}{a}\right)\left(\frac{1}{N}\right)\left[\sum_{j=1}^{i-1} \sum_{k=j+1}^{i} c_k - \sum_{j=i+1}^{N} \sum_{k=i+1}^{j} c_k \right] \qquad (5.9)$$

Substitution of the optimal demand levels into the fare product demand functions (equation 5.1) yields the optimal price levels for the OD market :

$$P_i^* = P_0 - a\left(\frac{i}{N}\right)\text{Cap} + \left(\frac{1}{N}\right)\sum_{j=1}^{i}\left[\sum_{m=1}^{j-1} \sum_{k=m+1}^{j} c_k - \sum_{m=j+1}^{N} \sum_{k=j+1}^{m} c_k \right] \qquad (5.10)$$

Finally, for the constant cost model with linear demands, the Lagrange multiplier term has the following form at optimality :

$$\lambda^* = P_0 - \left(\frac{N+1}{N}\right)a\text{Cap} - \left(\frac{1}{N}\right)\sum_{j=1}^{N} (N\text{-}j+1)\, c_j \qquad (5.11)$$

when the capacity constraint is binding.

The constrained optimal solution is valid in cases where $\sum_{i=1}^{N} Q_i \geq Cap$ for the aircraft serving the OD market.

5.2.3 Numerical Demonstration of Base Case Linear Constant Cost Model

Performance of numerical analysis using the constant cost model with linear demands is useful to demonstrate the testable applications. A three fare product "base case" is used for the analysis with the following OD market parameters :

1) P_0 = 600
2) a = 3
3) Cap = 150
4) c_1 = 0
5) c_2 = 75
6) c_3 = 25

Consequently, the linear demand functions for each of the three fare products available on the single OD market flight are :

$$P_1 = 600 - 3Q_1 \tag{5.12}$$

$$P_2 = 600 - 3(Q_1+Q_2) - 75 \tag{5.13}$$

$$P_3 = 600 - 3(Q_1+Q_2+Q_3) - 100 \tag{5.14}$$

The demand functions place the maximum willingness to pay of any passenger in the market for fare products 1, 2, and 3 at \$600, \$525, and \$500, respectively. Moreover, the airline must lower the fare by three dollars to attract an additional passenger into booking any of the fare products.

The market is served by a single 150 seat aircraft. In the base case, passengers are assumed to be perfectly segmented between the unrestricted fare product 1 and the more-restricted fare products 2 and 3. Fare product 1 is unrestricted and the

192

cost associated with accepting fare product 2 or 3 instead of the fare product 1 is $75 or $100, respectively. The degradation cost associated with the unrestricted fare product is, of course, zero. Fare product 2 has greater similarity (in terms of imposed consumer costs) to fare product 3 than to the unrestricted product. The base case represents a situation in which the lower-priced fare products are much less desirable than the unrestricted, similar to the "value pricing" plan initiated by American Airlines in 1992 which offered seven and fourteen day advance purchase, non-refundable excursion fares in addition to the unrestricted product.

The above OD market characterization is used as the basis against which all variations of the model are tested. Neither passenger diversion nor booking limits are considered in the base case. In other words, the following parameters apply in the base case formulation :

1) d_{12} = 0
2) d_{13} = 0
3) d_{23} = 0
4) BL_1 = ∞
5) BL_2 = ∞
6) BL_3 = ∞

The above parameters are varied to measure their effects on the optimal solutions later in the chapter when passenger diversion and booking limits are formally addressed. A derivation of the base case example having variations of a single parameter, are named, for example, "base case varying d_{13}" when the diversion percentage from fare product 1 to fare product 3 is changed.

The optimal values of price level, passenger demand, and airline revenue by fare product appear below in Table 5.1 for the base case :

Product	Passengers	Price	Revenue
FP1	64.58	$406	$26235
FP2	39.58	$213	$8410
FP3	31.25	$94	$2930

Table 5.1 Three Fare Product Base Case

The optimal solution would have a binding capacity constraint at approximately 135 seats with the Lagrange multiplier term associated with the capacity constraint near zero (*e.g.* an additional seat does not observably change the optimal solution or decision variables when this market is served by a 135 seat aircraft).

5.2.4 Varying the Aircraft Capacity

The results of the analysis vary with the capacity of the aircraft serving the market. Variations in the capacity constraint, however, are only interesting in the case of joint price level optimization when the capacity constraint is binding. A single unit relaxation of a binding capacity constraint usually results in the acceptance of an additional passenger in the lowest fare class except when the shadow price associated with the capacity constraint (*e.g.* Lagrange multiplier term) is zero. While varying the capacity by a single seat only changes the seat allocation by one seat, when using the joint price level optimization, the optimal price levels and passenger mix may shift as well.

It is evident from equation 5.10 that an increase in the capacity serving the market will unambiguously decrease all optimal fare product price levels. Figure 5.1 shows the decrease in all optimal fare product price levels occurring in the base case varying capacity until the joint optimization capacity constraint ceases to be binding at 135 seats. The number of passengers carried, on the other hand, increases up to the point where the capacity constraint no longer binds the solution as seen in Figure 5.2. Simple optimization sensitivity analysis guarantees that relaxation of a strongly binding constraint increases the objective function. The Lagrange multiplier term presented in Equation 5.11 represents the marginal value of an additional unit of capacity (*e.g.* another seat) in the OD

194

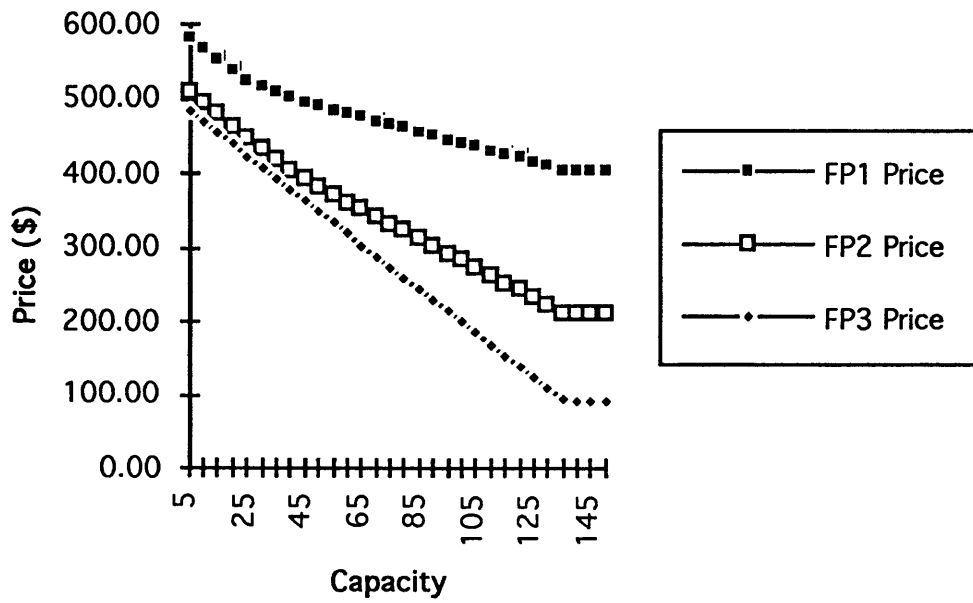## Base Case Varying Capacity - Prices by Fare Product

Figure 5.1 - Base Case Varying Capacity - Prices by Fare Product

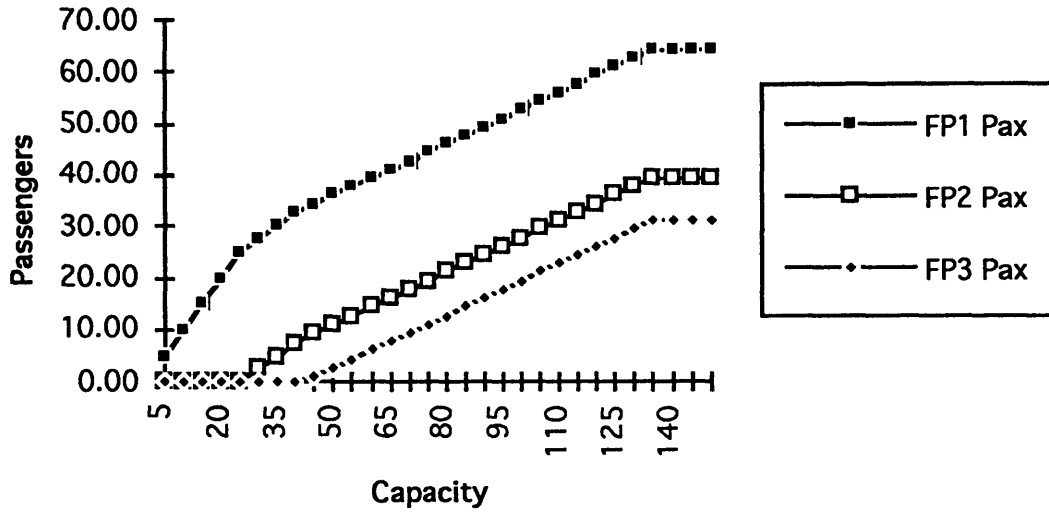**Base Case Varying Capacity - Passengers by Fare Product**

Figure 5.2 - Base Case Varying Capacity - Passengers by Fare Product

market. Thus, a revenue improvement is guaranteed with increased capacity in all capacity constrained cases except where the capacity constraint is loosely binding (e.g. the Lagrange multiplier term equals zero, making the revenue effect of increased capacity neutral). Figure 5.3 shows the increasing revenues in the OD market resulting from relaxations of the capacity constraint up to the unconstrained optimal total demand level where the objective levels off.

Figure 5.4 shows the effect of increased capacity on the traffic mix in the OD market. The optimal traffic mix is dominated by the least-restricted fare product for low capacities relative to demand. This is expected since when the number of seats available are quite limited, it benefits the airline to charge higher prices (as seen in Figure 5.1) and sell a greater proportion of those seats to the passengers with higher values of willingness to pay. For greater capacities, the unrestricted fare product progressively loses dominance until the capacity constraint ceases to be binding. Even so, the unrestricted product still maintains approximately fifty percent of all passenger traffic under perfect segmentation in the optimization base case varying capacity.

5.2.5 Varying the Degradation Costs Facing Consumers

The imperfection of second degree price discrimination techniques can be addressed in two ways, through the increased costs imposed on passengers by purchase restrictions and the inability of purchase restrictions to identify and segment the population as intended by the airline. For the former problem, varying the degradation costs associated with accepting increased purchase restrictions can be used to measure the effects of imperfect segmentation techniques. The effects of varying the degradation costs are addressed here. The latter problem can be addressed through different relaxations of the passenger segmentation assumptions (e.g. incorporating passenger diversion) addressed in the next section.

The unconstrained optimal price levels can be rewritten to show the effects of changing the different types of degradation costs. The transformed optimal price levels isolate the effects of more-restricted and less-restricted fare products :

197

Base Case Varying Capacity - Total Revenue
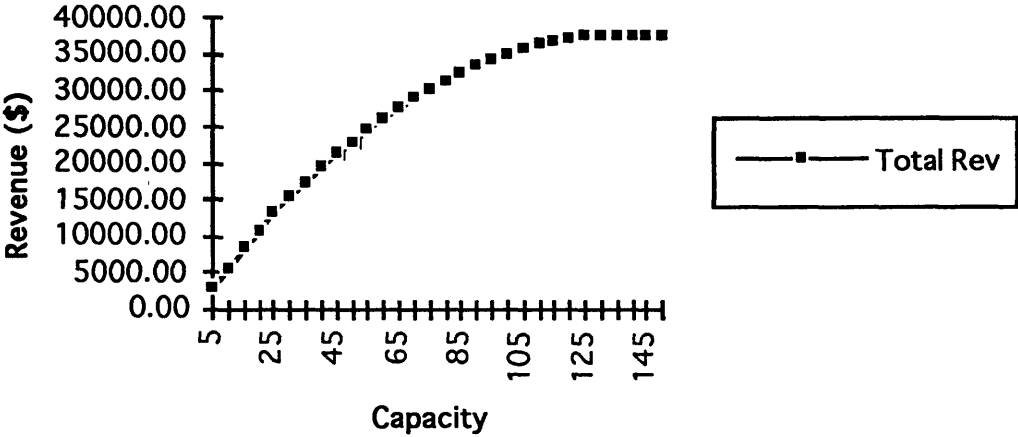
Figure 5.3 - Base Case Varying Capacity - Total Revenue

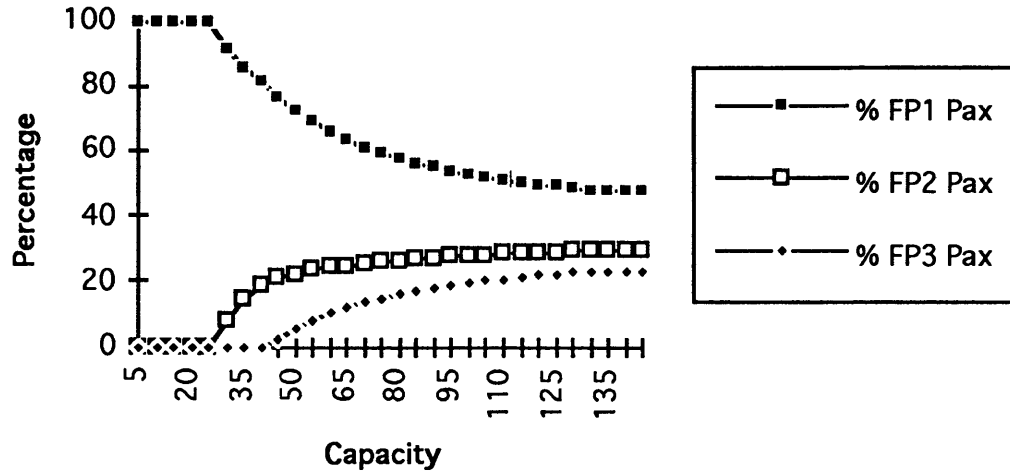**Base Case Varying Capacity - Percent of
Traffic by Fare Product**

Figure 5.4 - Base Case Varying Capacity - Percent of Traffic by Fare Product

$$P_i^* = \left[1 - \left(\frac{i}{N+1}\right)\right]P_0 - \left[1 - \left(\frac{i}{N+1}\right)\right]\sum_{k=1}^{i} kc_k + \left(\frac{i}{N+1}\right)\sum_{k=i+1}^{N} kc_k - \sum_{k=i+1}^{N} ic_k \quad (5.15)$$

First, the effects of the less-restricted fare products are shown. The effects of the degradation costs associated with the fare products less restricted than i on the optimal price level P_i are completely contained within the second of the four terms in equation 5.15. The following bound clearly holds for each (positive) fare product i and any number of fare products N :

$$0 \leq \left(\frac{i}{N+1}\right) < 1 \quad (5.16)$$

The effect of the second term of equation 5.15 on optimal price levels is thus, unambiguously negative. As the degradation costs associated with the less-restricted fare products increase, the optimal price level decreases. Optimal price levels vary inversely with the degradation costs associated with the less-restricted fare products. Therefore, airlines may achieve increases in the optimal price levels of their fare products through the reduction of the degradation costs associated with the higher-priced fare products.

Next, the effects of the more-restricted fare product degradation costs are explored. The degradation costs associated with the fare products that are more restricted than fare product i are contained in both the third and fourth terms of equation 5.15. It is easily shown that the following inequality holds in all cases :

$$\left(\frac{i}{N+1}\right)\sum_{k=i+1}^{N} kc_k < \sum_{k=i+1}^{N} ic_k \quad (5.17)$$

The above inequality reveals that as the degradation costs associated with the more-restricted fare products increase, the optimal price levels of any fare product fall. In other words, the third term of equation 5.15 dominates the fourth. The final result is that the negative component of more-restricted fare product degradation costs always exceeds the positive, resulting in an inverse proportionality between degradation costs and optimal price levels. Airlines may achieve increases in the optimal price levels for their fare products by the reduction of the degradation costs associated with lower-priced fare products.

In summary, a decrease in the degradation costs associated with any fare product results in an increase in the optimal price levels for the fare product in the unconstrained case. Simply stated, the constant cost model with linear demands tells us that as degradation costs decrease, optimal price levels rise in the absence of a binding capacity constraint. This result holds whether the decreases in degradation costs occur for the more-restricted or less-restricted fare products.

Transforming the solution (in terms of price levels) to the capacity constrained optimization identifies the effects of different types of degradation costs on optimal price levels in the other supply scenario :

$$P_i^* = P_0 - a\left(\frac{i}{N}\right)Cap + \left(\frac{1}{N}\right)\sum_{j=1}^{i}\sum_{m=1}^{j-1}\sum_{k=m+1}^{j} c_k$$
$$- \left(\frac{1}{N}\right)\sum_{j=1}^{i}\sum_{m=j+1}^{N}\sum_{k=j+1}^{i} c_k - \left(\frac{1}{N}\right)\sum_{j=1}^{i}\sum_{m=j+1}^{N}\sum_{k=i+1}^{m} c_k$$

$$(5.18)$$

Equation 5.18 shows that the optimal price level for each fare product i are a function of the degradation costs associated with all fare products. All of the degradation costs associated with the more-restricted fare products are contained within the fifth (and final) term on the right hand side of equation 5.18. Clearly, all costs contained in the final term have a negative impact on the optimal price level of fare product i. Thus, as in the capacity unconstrained case, the optimal price levels decrease as the degradation costs associated with the more-restricted fare products increase.

The effects on the optimal price level of fare product i of the degradation costs associated with fare products less restricted than i, on the other hand, are not determinate. For each fare product i, the effect on the optimal price level P_i of the degradation costs associated with less-restricted fare products (than i) has two components that impact the optimal price level differently. For any fare product j less restricted than i, the degradation costs associated with each fare product that is less restricted than j will lower the optimal number of passengers purchasing fare product j (as is evident from equation 5.9). Equation 5.1 demonstrates that a decrease in the number of passengers purchasing fare product j will, in turn, increase the optimal price level of fare product i. This

effect on optimal price levels is contained entirely within the third term on the right hand side in equation 5.18.

Using similar reasoning, the degradation costs associated with each fare product that is more restricted than each fare product j (less restricted than i) increases the optimal number of passengers purchasing fare product j and, thus, decreases the optimal price level for fare product i. The fourth term on the right hand side in equation 5.18 isolates this negative effect on the optimal price level P_i. The final result of the effect of the degradation costs of fare products less restricted than i on the optimal price level P_i is indeterminate because the size of each degradation cost is variable and neither of the two effects dominate.

Of greater concern to the airline is the effect on revenue of changes in the degradation costs. The total revenue for the market is simply :

$$\text{Total Revenue} = \sum_{i=1}^{N} P_i^* Q_i^* \qquad (5.19)$$

Since the quantities of passengers decrease and the optimal price levels increase with increases in the degradation costs in most cases, a closer look at the effect on revenue is required. An increase in the degradation costs associated with accepting a more-restricted fare product is equivalent to an inward shift of the demand function for that and all more-restricted fare products. As a result, optimal revenues never increase for the case in which demand for the fare products is lower. In other words, the decreased demand resulting from the increased passenger degradation costs more than counteracts the revenue increases of the price increases in all cases. Thus, high degradation costs are never in the interest of the airline and every attempt should be made to develop passenger segmentation techniques with the lowest possible degradation costs *ceteris paribus*.

The effect of increasing the degradation costs associated with a single fare product i is a decrease in revenue for the airline of $\Delta c_i Q_i$ when no reoptimization is performed. A reoptimization would result in a lower revenue (or at best equal) than in the case with less severe degradation costs because of the

decreased level of demand that results from the cost increases. The actual
revenue result depends upon the size of the new optimal price and traffic levels.

5.2.6 Variations of the OD Market Demand Levels

The variations of the demand function parameters in the constant cost model
with linear demands, in general, provide intuitive results. An increase in the
parameter P_0, or the maximum willingness to pay of any passenger, results in
higher demand levels and higher optimal price levels for both the capacity
constrained and unconstrained cases. Mathematically, these results are seen in
equations 5.4, 5.9, 5.5, and 5.10, respectively. Clearly, the optimal revenue also
increases with increased willingness to pay in the travel population.

Increases in the slope of the demand functions produce a downward effect on
optimal price levels for both the capacity constrained and unconstrained cases.
The effect on passenger demand of changes in the slope, however, differ in the
capacity constrained and unconstrained scenarios. Passenger demand at
optimality decreases with increases in the demand function slope in the
unconstrained case. Accordingly, the effect on revenue in the unconstrained case
is downward for any increases in the slope. This is also intuitively pleasing since
an increase in demand function slope represents increased sensitivity to price
level within the travel population.

The effect on passenger demand of changes in the slope of the demand function
in the capacity constrained case is not determined due to the variability of the
degradation costs. If the degradation costs associated with fare products less
restricted than i are large relative to the costs accompanying the more-restricted
fare products, then the slope has an upward effect on the optimal Q_i. Large
degradation costs on the fare products more restricted than i produce the
opposite effect on Q_i with respect to increases in demand function slope.
Unfortunately, the effect on revenue is not determined either.

5.2.7 Benefits of Introducing More Fare Products

As the number of differentiated fare products offered in the market increases, the
revenue potential for the airline also increases. Consequently, the ability to

203

perform an increasingly disaggregate passenger segmentation enhances the expected revenue of the airline. Unfortunately for the airline, there are diminishing returns to increasing the number of fare products offered on the market, even assuming perfect passenger identification and segmentation . The constant cost model demonstrates that the revenue improvement per fare product becomes progressively smaller as more fare products are introduced. In the limit, the revenue contribution of adding fare products approaches zero.

Lower levels of passenger demand mitigate the revenue impact of introducing additional fare products. Recall that increasing the degradation costs associated with any of the fare products is equivalent to an inward shift in the demand function for that product. The resulting optimal revenue after a demand function has shifted inward cannot be greater than that prior to the shift. Thus, the generalized cost model in the absence of degradation costs determines an upper bound on the effects on total revenue of the introduction of additional fare products onto the market.

The generalized cost model improves upon the existing models of airline fare product differentiation by introducing the demand curbing effects of perceived passenger degradation costs associated with the more-restricted fare products. Under the existing models (equivalent to the generalized cost model without degradation costs), airlines systematically overpredict the effects of introducing additional fare products to the market. While the diminishing returns on additional fare product introductions are still apparent in the previous models, their impact may be significantly less than predicted leading to unrealistic performance expectations. Nonetheless, an exploration of the upper bound on the revenue impacts of introducing additional fare products to the market is instructive and appears next.

For the capacity unconstrained case with no degradation costs, the optimal price levels and expected demand are shown below :

$$P_i^* = P_0 - \left(\frac{i}{N+1}\right)P_0 \tag{5.20}$$

$$Q_i^* = \frac{P_0}{(N+1)a} \tag{5.21}$$

204

Equation 5.20 demonstrates that the optimal price levels increase for each of the fare products as N increases. Equation 5.21 ensures a decrease in the demand for each fare product as additional fare products are introduced. In addition, an even distribution of passengers among the fare products is guaranteed by equation 5.21 with total demand for the flight of :

$$Q_{Total}^* = \frac{NP_0}{(N+1)a} \tag{5.22}$$

Total demand increases in N and approaches the asymptotic limit of :

$$\lim_{as\ N \to \infty} \left[Q_{Total}^* \right] = \frac{P_0}{a} \tag{5.23}$$

The total expected revenue in the optimal capacity unconstrained case for N fare products is :

$$Revenue_N^* = \left(\frac{N}{N+1} \right) \left[\frac{P_0^2}{2a} \right] \tag{5.24}$$

which shows the total OD market revenue to be increasing in N and approaching the asymptotic limit :

$$\lim_{as\ N \to \infty} \left(Revenue_N^* \right) = \frac{P_0^2}{2a} \tag{5.25}$$

The change in revenue experienced by the airline as the number of fare products offered is increased from N to N+1 for the capacity unconstrained case appears below :

$$\frac{\Delta Revenue^*}{\Delta N} = \left[\left(\frac{N+1}{N+2} \right) - \left(\frac{N}{N+1} \right) \right] \left[\frac{P_0^2}{2a} \right] \tag{5.26}$$

The above relation shows that the final result of the optimal price level increases among an increasing number of evenly distributed passengers is an increase in revenue for the market as additional fare products are introduced. Equation 5.26 also shows that the exact size of the revenue increase resulting from an additional fare product introduction depends not only upon the number of fare

products already offered on the market, but also on the demand parameters for the market. Although a revenue increase is guaranteed from the introduction of new fare products, in the limit, the revenue increment is driven to zero :

$$\lim_{\text{as } N \to \infty} \left(\frac{\Delta \text{Revenue}^*}{\Delta N} \right) = 0 \tag{5.27}$$

demonstrating the diminishing returns to additional fare product introduction in the optimal capacity unconstrained case.

For the capacity constrained case, the optimal price levels and expected demand are :

$$P_i^* = P_0 - a\left(\frac{i}{N}\right)\text{Cap} \tag{5.28}$$

$$Q_i^* = \frac{\text{Cap}}{N} \tag{5.29}$$

In contrast to the capacity unconstrained case, the total passengers cannot exceed the aircraft cabin capacity leaving total passengers constant and equal to capacity. Therefore, the increase in optimal price level that occurs for each fare product as N increases also results in increased revenues for the airline since passengers are evenly distributed among the fare products once again. The total revenue for the OD market flight is :

$$\text{Revenue}_N^* = P_0\text{Cap} - \left(\frac{N+1}{N}\right)\left[\frac{a\text{Cap}^2}{2}\right] \tag{5.30}$$

and approaches an asymptotic limit of :

$$\lim_{\text{as } N \to \infty} \left(\text{Revenue}_N^*\right) = P_0\text{Cap} - \left[\frac{a\text{Cap}^2}{2}\right] \tag{5.31}$$

In the capacity constrained case, the increase in total revenue received by the airline from offering an additional fare product (under perfect segmentation) is :

$$\frac{\Delta \text{Revenue}^*}{\Delta N} = \left[\left(\frac{N+1}{N}\right) - \left(\frac{N+2}{N+1}\right)\right]\left[\frac{aCap^2}{2}\right] \tag{5.32}$$

as the number of fare products is increased from N to N+1. As the number of fare products N becomes large, the revenue increment resulting from the introduction of an additional fare product becomes smaller. In fact, the revenue increment is driven to zero as N becomes very large :

$$\lim_{\text{as } N \to \infty} \left(\frac{\Delta \text{Revenue}^*}{\Delta N}\right) = 0 \tag{5.33}$$

Thus, the diminishing returns to fare product additions are replicated for the optimal capacity constrained case.

An alternative view of revenue potential is offered by examining total system welfare. The upper bound on total revenue achievable by the airline is equal to total welfare in the capacity unconstrained case. With a binding capacity constraint, the upper bound on total revenue equals the total welfare available to the passengers with the highest values of willingness to pay (up to capacity). The limit of total revenue as N approaches infinity is :

$$\text{Revenue UB} = \int_0^{Q_{max}} [P_0 - aq]dq \tag{5.34}$$

where
Q_{max} = the maximum demand allowable for the OD market flight

In the optimal capacity unconstrained case, equation 5.34 is equivalent to equation 5.25 since Q_{max} is equal to the maximum level of demand in the market (at zero price level) given in equation 5.23. Q_{max} is equal to the total number of available seats in the optimal capacity constrained case, making the solution of equation 5.34 equivalent to equation 5.31.

Some simple examples provide insight to the impact of the addition of fare products to the market. With no binding capacity constraint, 80% of the total

potential revenue gain can be achieved by offering four fare products (dividing equation 5.24 by equation 5.25). Ninety percent of the revenue can be guaranteed from the introduction of nine fare products in the unconstrained case. Introducing the next 90 fare products leads to only a 9% improvement in revenue (assuming perfect segmentation). The incremental revenue returns to additional fare products decrease quite rapidly as Figure 5.5 illustrates for the capacity unconstrained case. In the capacity constrained case, the total potential revenue gain is dependent upon the parameters P_0, a, and Cap (dividing equation 5.30 by equation 5.31). Hence, the exact revenue potential in the capacity constrained case is OD market specific.

Practical constraints also limit the ability of the airline to offer more fare products and should be considered when contemplating increased passenger segmentation attempts. When an airline introduces a new fare product, costs are incurred as a result of such factors as increased advertisement and added service offering complexity. The costs of offering additional fare products are likely to exceed the revenue benefit when several fare products are already offered. In addition, the ability to segment passengers into smaller populations becomes less likely when significant segmentation has already been achieved. The number of fare products that can be offered by a carrier is also limited by the capabilities of the computer reservations systems. The reservation system data bases are limited in the number of fare products that they store due to the shortcomings of the storage capacity of the reservations systems at many carriers. Finally, many consumers already resist the "complicated" fare product structure on the market today. An increase in the number of fare products may confuse or upset such consumers.

5.2.8 Effects of Degradation Costs on Fare Product Introductions

When degradation costs are included in the generalized cost model, the positive revenue impacts of introducing additional fare products to the OD market are mitigated. The generalized cost model with degradation costs is compared to the case of zero degradation costs in an exploration of the effects of perceived passenger cost on the revenue improvements available from increasing the number of fare products.

208

Varying the Number of Fare Products - Percent of Revenue Potential
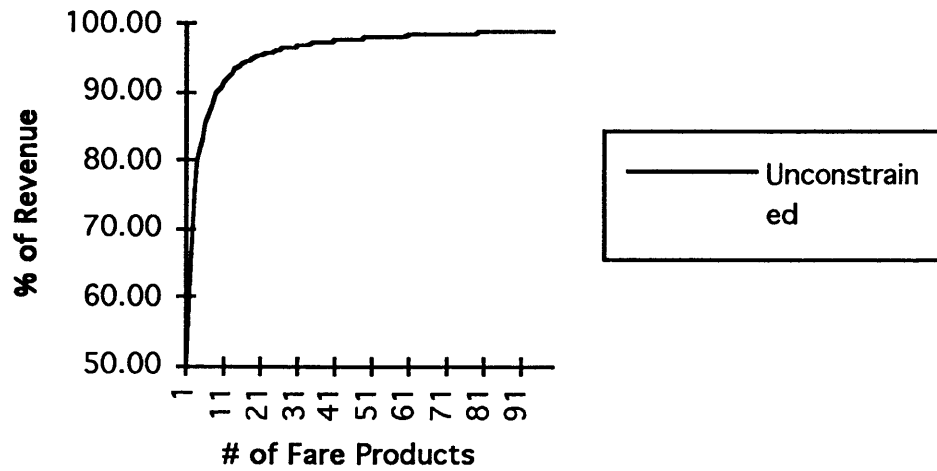
Figure 5.5 - Varying the Number of Fare Products - Percent of Revenue Potential

The degradation costs associated with each fare product on the market are set to $5 and then $10 in the examples presented. In other words, the (incremental) degradation cost c_i associated with each fare product i is equal to $5 and $10, respectively, in the two cases presented. For example, the total degradation costs associated with fare product 3 are $10 and $20, respectively, in the cases. Similarly, the fare product 4 degradation costs are $15 and $30. As expected, increasing degradation costs (c_i) reduce total optimal revenue increasingly as shown in Figure 5.6. The lower the perceived value of the fare products on the market, the less impact that additional fare product introductions have on total revenue.

Although lowered, the revenue impact remains non-negative when degradation costs are varied *ceteris paribus*. A non-negative revenue impact is guaranteed in all cases where the fare product structure is only varied by the addition of a single fare product (regardless of the associated degradation costs). Thus, although increased degradation costs mitigate the positive revenue impacts of fare product introductions onto the market, even large degradation costs cannot cause additional fare product introductions to have a negative revenue impact for the case of perfect passenger segmentation and identification.

5.3 Constant Cost Model with Linear Demands Allowing Passenger Diversion

The assumption that airlines are able to perfectly identify and segment the consumer population by their willingness to pay is relaxed in this section. The impacts of ineffective fences on the OD market can be quantified using the generalized cost model by varying the assumptions surrounding the passenger segmentation made by the airline. Unfortunately, obtaining analytical solutions to the constant cost model allowing passenger diversion becomes too complicated to yield intelligible results even for simple cases. As a result, only numerical analysis using the GAMS optimization package is performed when measuring the effects of passenger diversion. The optimal solutions to the model extensions with diversion are compared to each other and referenced to the no diversion case to make insights concerning the impacts of passenger diversion.

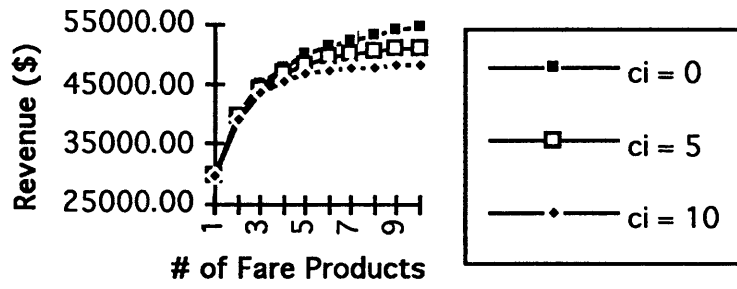**Base Case** Varying Number of
Fare Products - Total Revenue

Figure 5.6 - Base Case Varying Number of Fare Products - Total Revenue

Examples of three fare product markets are used to demonstrate the effects of passenger diversion on the optimal solution to the joint price level optimization problem through comparisons to the base case model (without diversion) and cross-comparisons varying input parameters. In the generalized cost model, diverting passengers are assumed to be initially segmented into purchasing a higher-priced, less-restricted fare product. Then, as a result of their ability to meet the imposed restrictions, they purchase a lower-priced fare product. Hence, passenger diversion may only occur from higher-priced, less-restricted fare products to lower-priced ones. In the initial three fare product example, fare product 1 consumers are allowed to divert to fare product 2 exclusively while fare product 2 consumers may not divert. This scenario is referred to as the base case varying d_{12}. Passengers are allowed to divert between the other fare products in subsequent examples.

It should be noted that the diverting passengers are still subject to the degradation costs associated with the more-restricted products that they purchase. In the base case varying d_{12}, for example, passengers diverting from fare product 1 incur a cost of $75 when they purchase fare product 2. In other words, although the fencing mechanism is imperfect, the diverting passengers are still inconvenienced by the added travel restrictions. The sole benefit received by the diverting passenger is the added consumer surplus received from the lower price level of the more-restricted fare product (net of degradation costs).

5.3.1 Basic Impacts of Passenger Diversion

The passenger diversion effect of primary concern to airlines is the revenue dilution that results from passengers purchasing fare products less expensive than intended. When passengers divert to the lower-priced fare products as a result of their ability to meet the purchase restrictions imposed by airlines, the revenue received by a carrier may be reduced in addition to other optimality conditions being changed. The negative revenue impacts of diversion are expected to increase with the number of passengers diverting. Figure 5.7 shows the expected decline in revenue for the optimal price level solution as passenger diversion between fare product 1 and fare product 2 increases in the base case of the linear constant cost model varying d_{12}. In the figure, and all future figures
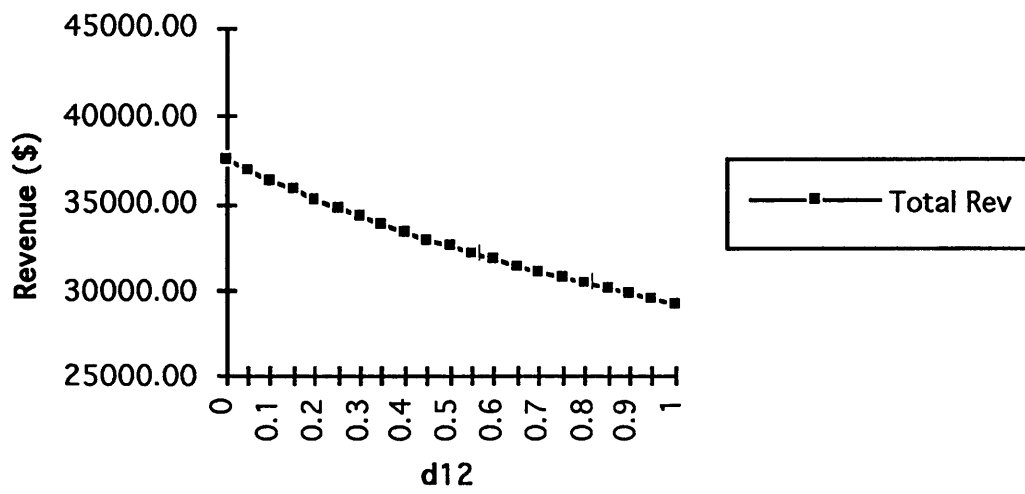
Base Case Varying d12 - Total Revenue

Figure 5.7 - Base Case Varying d12 - Total Revenue

addressing passenger diversion, d_{ij} is defined as the percentage of diversion occurring from fare product i to fare product j.

A decrease in optimal revenue does not, however, always result from increased diversion. In the case where diversion occurs from fare product 1 to fare product 3 exclusively (seen in Figure 5.8), the impact of increased diversion is revenue-neutral when a 60% diversion rate is exceeded. Thus, while increasing the level of passenger diversion generally results in a decrease in optimal airline revenues, the effect can be revenue neutral as witnessed in the specific case (base case varying d_{13}) presented in Figure 5.8.

Without diversion, the unrestricted fare product commands nearly seventy percent of the revenue for the entire flight. As the percentage of diversion from fare product 1 to fare product 2 increases, the amount of revenue received from the unrestricted fare product decreases and, in the limit, is driven to zero. Figure 5.9 highlights the declining importance of the unrestricted fare product to optimal revenue performance as fare product 1 passengers divert increasingly. In contrast, the revenue received from fare product 2 at optimality increases as a result of the increased demand from diverting passengers. In fact, fare product 2 possesses over seventy percent of the revenue share among the fare products at 100% diversion exceeding the revenue share dominance of fare product 1 in the no diversion base case.

Interestingly, not only is the revenue once earned by fare product 1 transferred to fare product 2 in the optimal solution, but also to fare product 3 as seen in Figure 5.10. There is upward pressure on both the optimal price levels and demand for fare product 3 as Figures 5.11 and 5.12 demonstrate. The percentage of revenue earned by fare product 3 with increased fare product 1 diversion increases from under 9% to almost 29% aided by an increase in traffic share from 23% to 46%. As expected, in the face of increasing fare product 1 diversion, fare products 2 and 3 increase in importance in both revenue and passenger share. Overall, the number of passengers carried for each of the lower-priced fare products increases while the total number of passengers carried decreases as shown in Figures 5.12 and 5.13, respectively. Therefore, while the importance of the lower-priced fare products increases (in terms of both traffic and revenue per fare product), the final effect on traffic and revenue is negative in the presence of diversion.

214

Base Case Varying d13 - Total Revenue
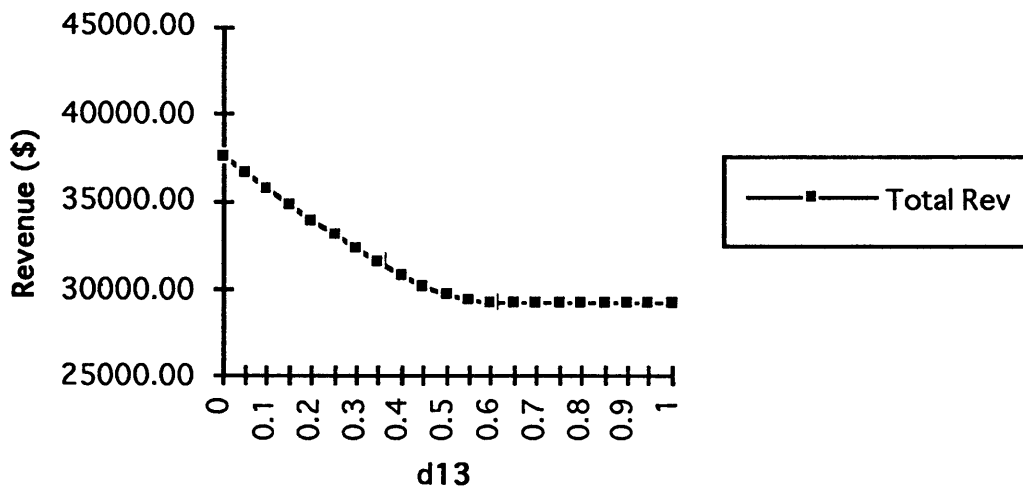
Figure 5.8 - Base Case Varying d_{13} - Total Revenue

Base Case Varying d12 - Percent of
Revenue by Fare Product
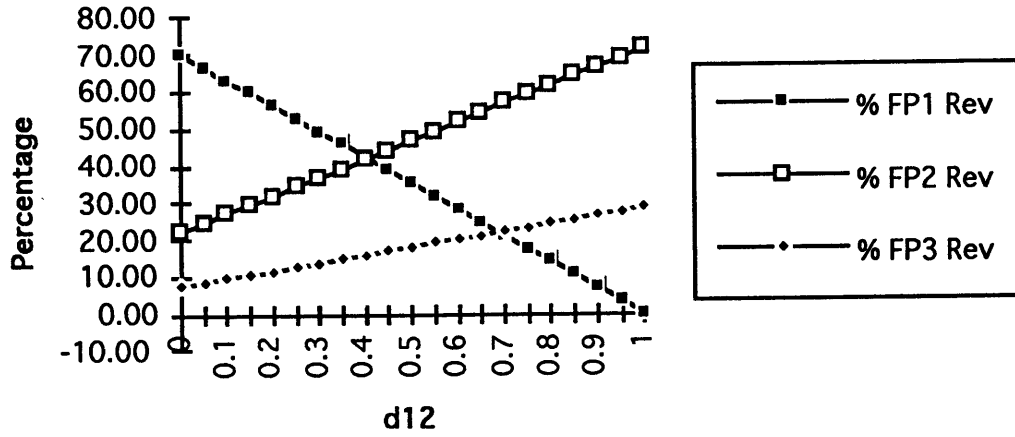
Figure 5.9 - Base Case Varying d12 - Percent of Revenue by Fare Product

Base Case Varying d12 - Revenue by Fare Product

Figure 5.10 - Base Case Varying d12 - Revenue by Fare Product

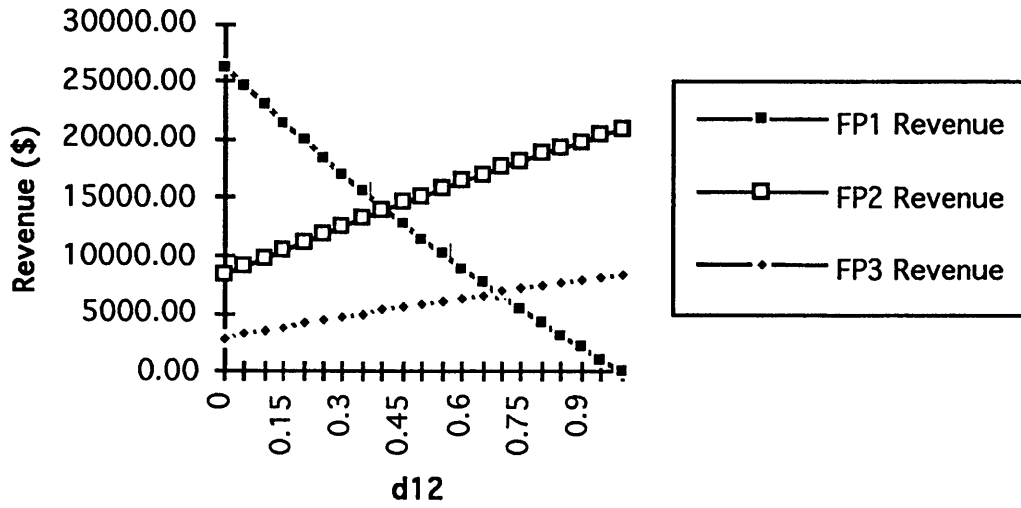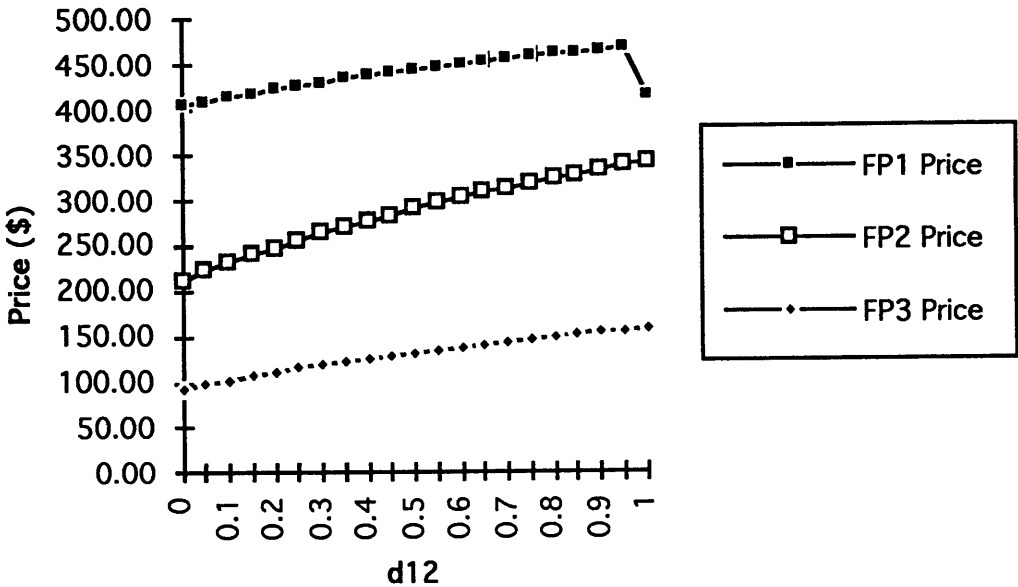Base Case Varying d12 - Prices by Fare Product

Figure 5.11 - Base Case Varying d12 - Prices by Fare Product
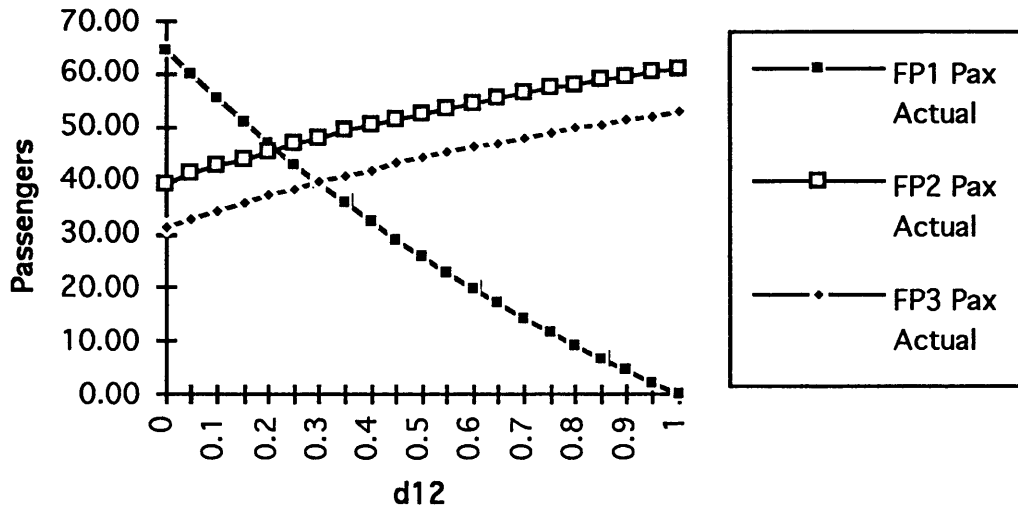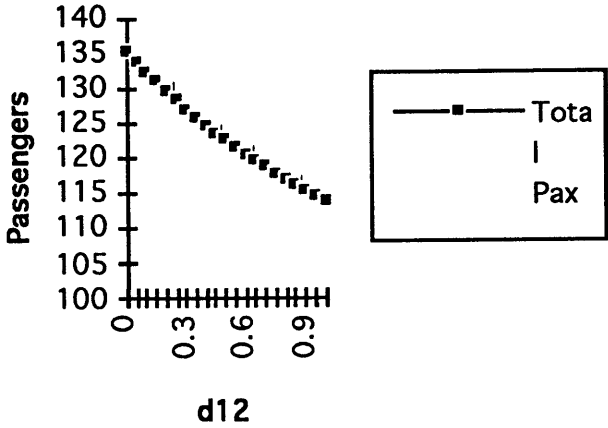
Base Case Varying d12 - Passengers by Fare Product

Figure 5.12 - Base Case Varying d12 - Passengers by Fare Product

Base Case Varying d12 -
Total Passengers

Figure 5.13 - Base Case Varying d$_{12}$ - Total Passengers

Looking more closely at the optimal price levels, Figure 5.11 shows that passenger diversion from fare product 1 to 2 puts upward pressure on the optimal fare levels for all three fare products. It is reasonable that an airline would want to accept fewer passengers at a higher fare level in the presence of a high diversion percentage for the unrestricted fare product. The results of the optimization considering passenger diversion suggest that a revenue maximizing airline should raise price levels to counteract the revenue dilution effects of diversion.

As the diversion percentage increases, the optimal solution to the three fare product example approximates the two fare product case. The only differences between the two fare product case and the three fare product case with 100% diversion for fare product 1 are the added degradation costs for the two "selling" fares. In fact, the optimal price levels from the joint price level optimization for two fare products with zero degradation costs on the restricted fare product yield results identical to the 3 fare product case with d_{12} diversion. With degradation costs, however, the optimal three fare product solution with 100% fare product 1 diversion produces less revenue than the optimal two fare product case. Therefore, two fare products with an effective fence between them may improve revenue for the airline over three fare products having significant diversion when degradation costs are considered. Airlines should pay close attention to diversion considerations when designing new fare product for the market because of the revenue dilution effects that may occur from higher-priced fare product passengers buying down.

The cost of diversion in terms of total revenue is expected to decrease as the fare products from which passengers divert become more restricted (and less expensive). In this way, when a passenger segmentation scheme fails for the passengers with the highest values of willingness to pay, the revenue impact is potentially greater than when passengers with lower values of willingness to pay divert. Figure 5.14 presents a comparison of the revenue impacts of passenger diversion in the cases where passengers divert from fare products 1 to 2 and 1 to 3. In the figure, for any fixed level of passenger diversion from fare product 1, the negative revenue impact is greater when passengers divert to fare product 3 than when diversion occurs to fare product 2.
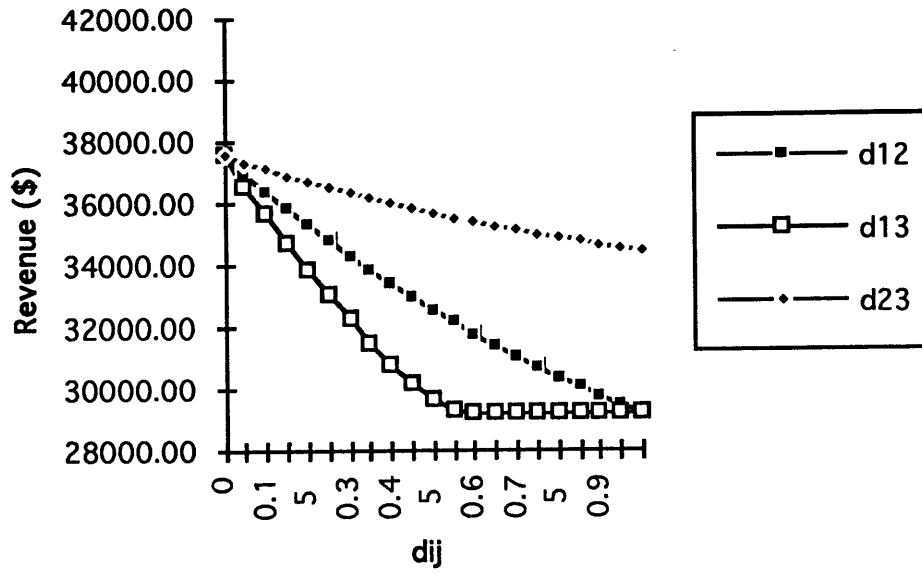
Figure 5.14 - Base Case Varying Diversion Type - Total Revenue

Similarly, passengers with lower levels of willingness to pay who divert are expected to have less of an impact on airline revenues. This effect is reflected by the comparison of passenger diversion from fare product 2 to 3 to that from fare product 1 to 3 also presented in Figure 5.14. The negative revenue effect of diversion from fare product 2 to fare product 3 is significantly less than that from fare product 1. Thus, the cost to the airline of diversion from passengers with higher values of willingness to pay who purchase less-restricted fare products than intended is greater *ceteris paribus*.

5.3.2 Varying the Degradation Costs Associated with Fare Products

The base case provides only a single instance of a fare product structure facing consumers in the OD market. Changes in the fare product structure or consumer perceptions of the costs associated with purchase restrictions can change the optimal solutions to the joint price level optimization problem. In this section, the sensitivity of the generalized cost model to changes in the fare product structure (as exhibited in the degradation costs) is tested. The results seek to demonstrate the effect on optimal price levels, passenger demand, airline revenue, and societal welfare of varying the fare product structure (degradation costs) when airline passenger segmentation schemes are imperfect.

The degradation costs associated with fare product 2 are varied from the base case model to examine the sensitivity of the optimal results to changes in the degradation costs. Specifically, the degradation cost associated with fare product 2 relative to fare product 1 varies from zero to $100 in increments of $25. Fare products 1 and 2 are identical in the eyes of consumers when the degradation cost associated with fare product 2 is zero. A degradation cost of $100, on the other hand, makes fare product 2 identical to fare product 3 since the total degradation cost associated with fare product 3 is held at a constant $100 more than the unrestricted fare product (1) for all of the tests. As before, the results of the joint price level optimization are calculated under varying passenger diversion (d_{12}) percentages.

As the degradation costs associated with the restricted fare products increase, the cost of diversion to the airline (and welfare) increases. This is expected since each diverting passenger able to purchase the more-restricted fare product incurs

the (increased) associated degradation costs. Figure 5.15 shows the decrease in total optimal revenues that occur as fare product 2 is made increasingly less attractive to consumers. The total market welfare reduction appears in Figure 5.16 and, as expected, increased degradation costs result in larger welfare reductions.

Recalling that an isolated increase in the degradation costs associated with a single fare product is equivalent to an inward shift in the demand for that fare product, decreased demand levels for fare product 2 are shown in Figure 5.17 as the fare product 2 degradation costs are increased. In contrast, the total number of passengers actually purchasing fare products 1 and 3 at optimality actually increase as shown in Figures 5.18 and 5.19, respectively. This makes sense since increasing fare product 2 degradation costs makes fare products 1 and 3 more attractive relative to fare product 2 *ceteris paribus*. Overall, however, the fare product 2 traffic losses dominate and the total passenger levels decrease with the increased degradation costs as Figure 5.20 demonstrates.

Figures 5.21 and 5.22 show that the optimal price levels associated with fare products 1 and 2, respectively, are higher with lower fare product 2 degradation costs. Fare product 2 looks relatively less attractive to fare product 1 as fare product 2 degradation costs increase. Fare product 1 commands a higher price level *ceteris paribus* as it becomes relatively more attractive to consumers. In addition, there is increasing pressure to lower the number of passengers diverting to the less-attractive fare product since fare product 2 is priced lower as c_2 increases. Lower fare product 1 demand can only be achieved by raising the its price level.

A discontinuity of the fare product 1 price levels appears in Figure 5.21 when the passenger diversion percentage between fare products 1 and 2 reaches 100%. Not only do the fare product 1 price levels drop significantly, but their "ordering" with respect to degradation costs also shifts. In the case with 100% diversion occurring from fare product 1 to 2 only, it should be noted that the fare product 1 price level no longer effects the objective function allowing alternate optima to exist for the joint price level optimization one of which is the situation appearing in Figure 5.21; another is the continuous solution. In fact, the optimal fare product 1 price level can take on any value greater than or equal to the

Base Case Varying Degradation Costs - Total Revenue
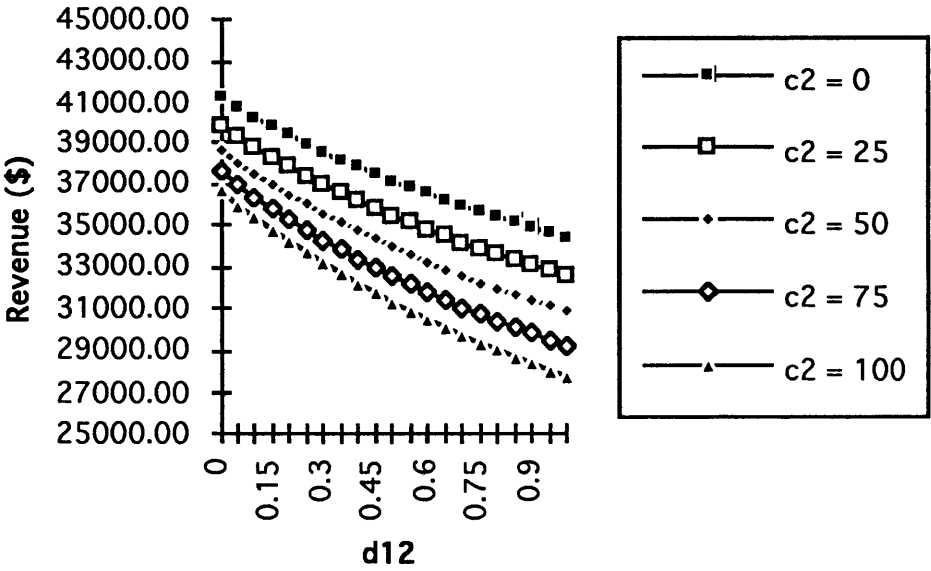
Figure 5.15 - Base Case Varying Degradation Costs - Total Revenue

Figure 5.16 - Base Case Varying Degradation Costs - Total Welfare

Figure 5.17 - Base Case Varying Degradation Costs - FP2 Passengers Carried

Base Case Varying Degradation Costs - FP1 Passengers Carried

Figure 5.18 - Base Case Varying Degradation Costs - FP1 Passengers Carried

**Base Case Varying Degradation Costs
- FP3 Passengers Carried**
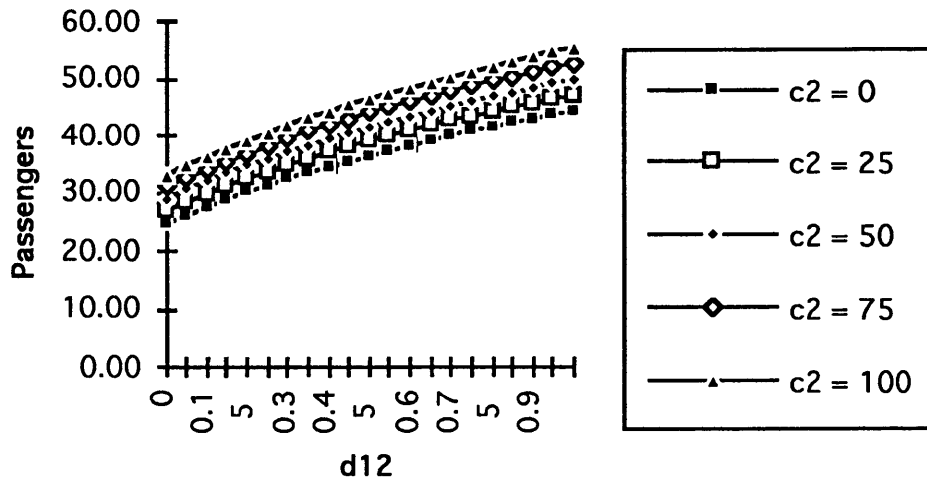
Figure 5.19 - Base Case Varying Degradation Costs - FP3 Passengers Carried
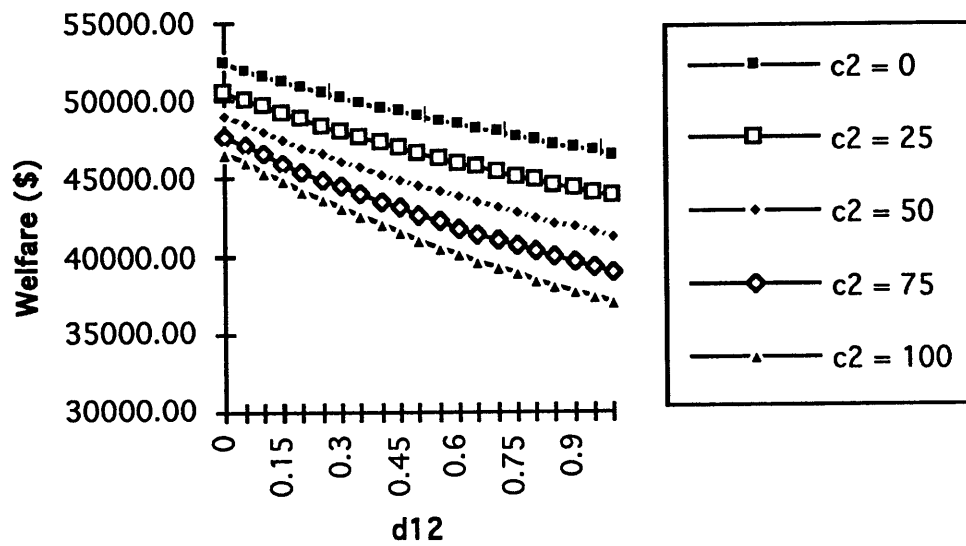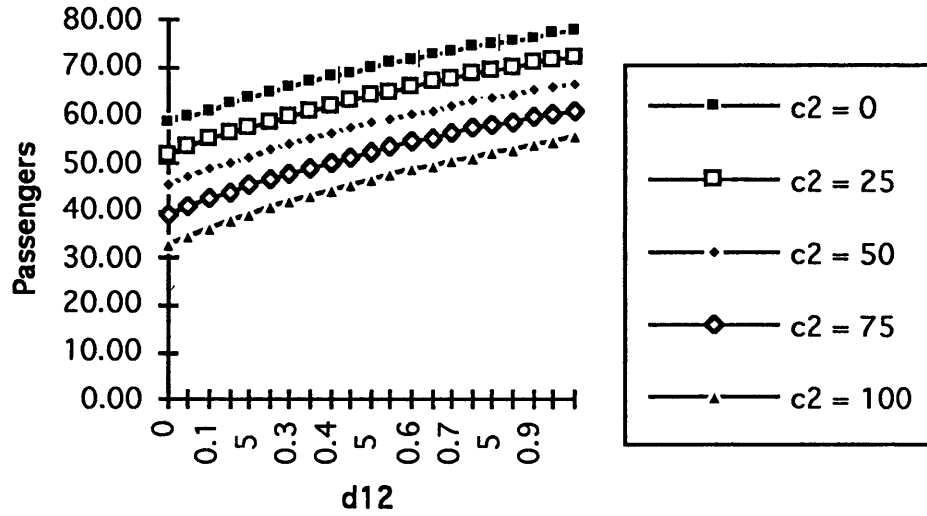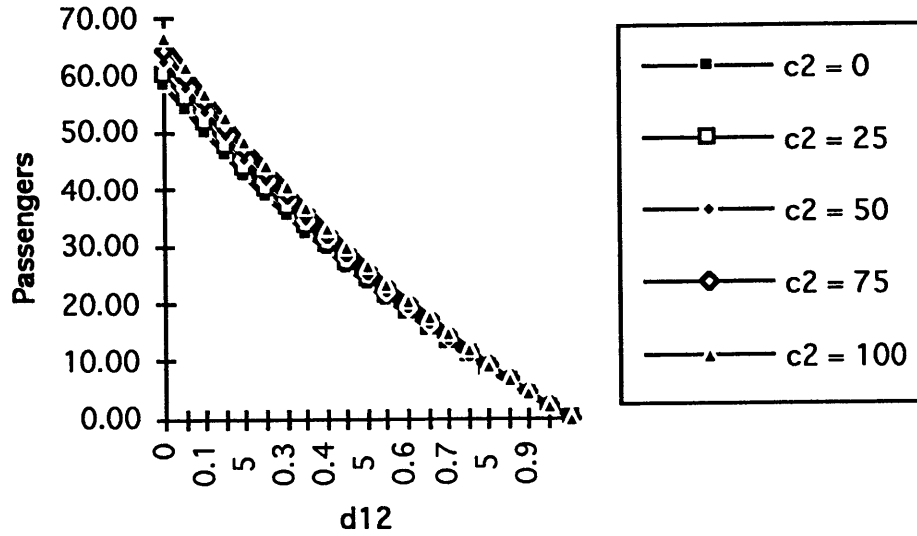
Base Case Varying Degradation Costs - Total Passengers
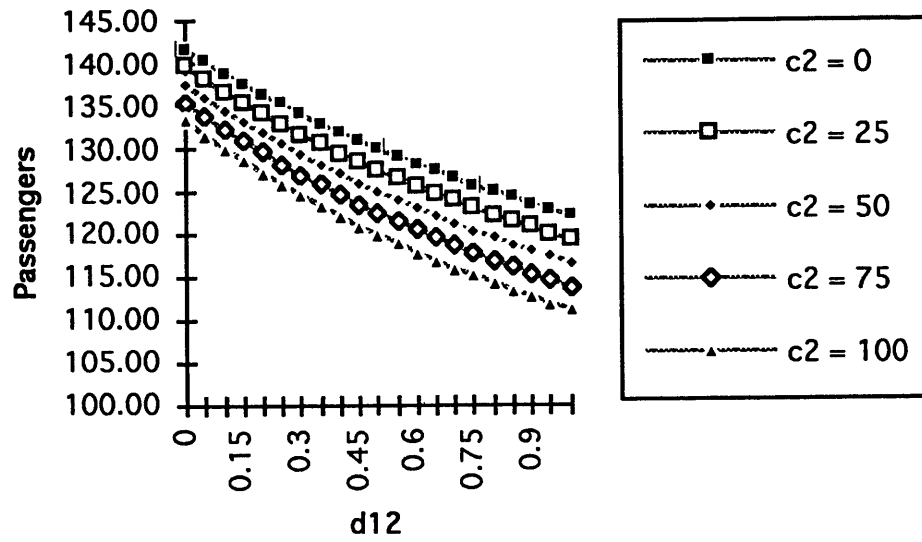
Figure 5.20 - Base Case Varying Degradation Costs - Total Passengers

Base Case Varying Degradation Costs - FP1 Price Level
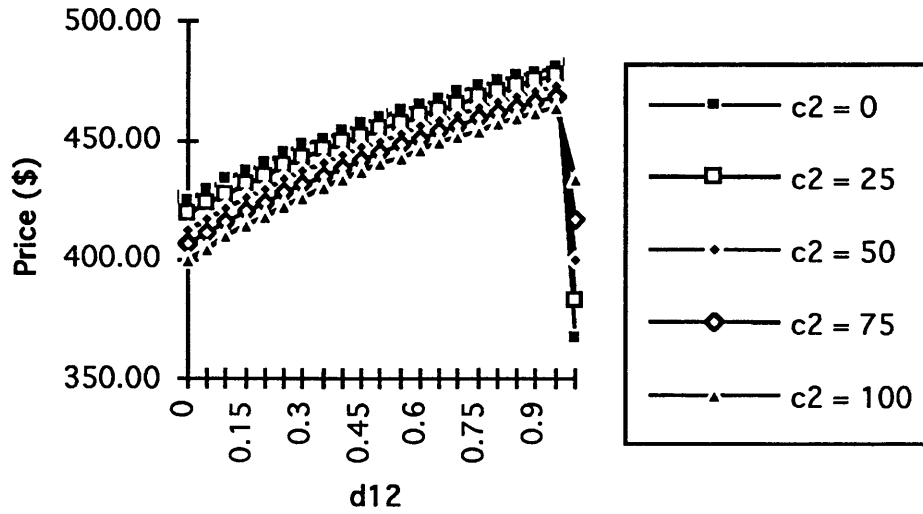
Figure 5.21 - Base Case Varying Degradation Costs - FP1 Price Level

Figure 5.22 - Base Case Varying Degradation Costs - FP2 Price Level

optimal fare product 2 price level. Therefore, the discontinuity shown in the figure need not exist.

The optimal fare product 2 price level has upward pressure with increased diversion since captive demand is being shifted down from fare product 1. The diverting fare product 1 passengers are captive to fare product 2 as long as P_2 remains less than $P_1 + c_2$. It is therefore in the interests of the airline to charge the diverting fare product 1 passengers the highest possible price level for fare product 2. Thus, the optimal fare product price 2 level must shift upward increasingly to compensate for the downward pressure applied by the increasing degradation costs. Optimal fare product 3 price levels, on the other hand, are higher for the fare product structures with higher fare product 2 degradation costs as Figure 5.23 demonstrates. The higher values of the optimal price level for fare product 3 reflect its relative improvement (in consumer perceptions) compared to fare product 2 and the lower number of total passengers carried when degradation costs are increased.

5.3.3 Tradeoff Between Degradation Costs and Passenger Diversion

Ideally, airlines would like to perfectly segment passengers by their willingness to pay while imposing no fare product degradation costs upon them. To this end, airlines attempt to design fare product purchase restrictions that provide the best segmentation while imposing the least cost on consumers. Due to the practical constraints of fare product design, however, airlines generally face a choice between the severity of the costs imposed by the segmentation device and the level of passenger diversion occurring. Once again, the airlines must confront the demand stimulation versus revenue dilution tradeoff. The nature of this tradeoff can be quantified using the assumed passenger diversion percentages and degradation costs input to the generalized cost modeling framework. The results of such an analysis can be used to measure the benefits and pitfalls of proposed fare product modifications aimed at preventing passenger diversion.

For the three fare product base case with diversion between fare products 1 and 2 only, imagine that the restrictions accompanying fare product 2 may be modified to reduce the degradation costs (associated with fare product 2) by \$50. To do so,

233

**Base Case Varying Degradation Costs
- FP3 Price Level**
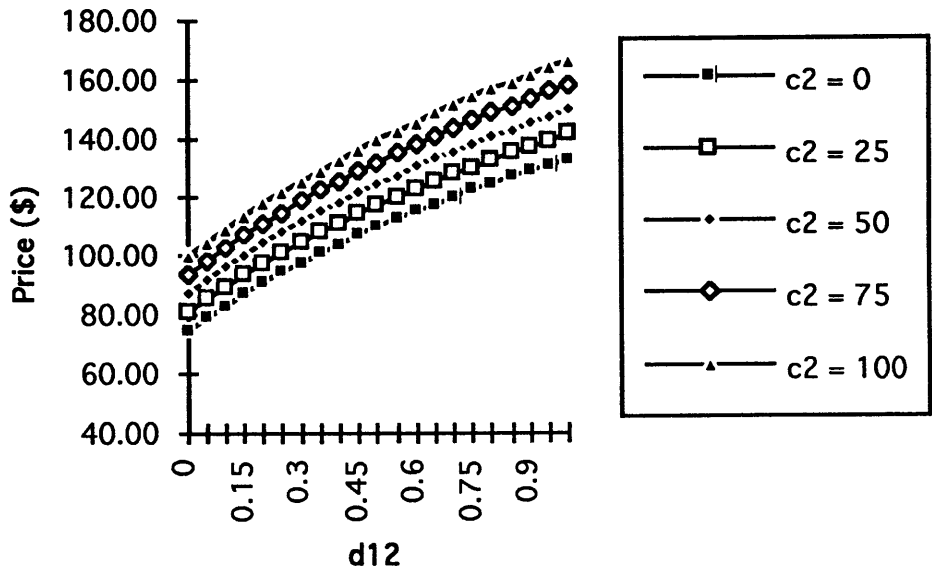
Figure 5.23 - Base Case Varying Degradation Costs - FP3 Price Level

however, results in an increase in the passenger diversion percentage from 20% to 40% as a result of the decreased effectiveness of the more lenient purchase restrictions. The airline faces a choice between the decrease in the degradation costs with an increase in the passenger diversion rate and maintaining the status quo.

Examining Figure 5.15, a $50 dollar decrease in the fare product 2 degradation costs and the associated increase in passenger diversion from 20 to 40% results in a net increase in expected airline revenues of $917 over the prevailing conditions. From a revenue perspective, the airline would be advised to modify the current fare product structure by reducing the degradation costs associated with fare product 2 and accepting the increase in passenger diversion.

Although fare products with higher degradation costs have more negative revenue impacts in the constant cost model without diversion *ceteris paribus*, the positive revenue effects of decreased passenger diversion may counteract the increased degradation costs. While airlines should strive to introduce fare products with the lowest possible degradation costs, this should not be done at the expense of large increases in diversion. Airlines considering new fare product designs must consider the effects of both perceived consumer degradation costs as well as diversion to best enhance their fare product structures.

5.3.4 Admitting that Passenger Diversion Exists May Improve Results

The cost to airlines resulting from passenger diversion can be substantial as shown in the above examples. Although it is unlikely that the airline can prevent passenger diversion from occurring, the revenue dilution effects of diversion can be minimized if the nature of passenger diversion is known. Simply stated, an airline that admits to having a diversion problem can recuperate some of its losses through improved pricing decisions. This subsection seeks to demonstrate that the costs of denying the existence of passenger diversion can be great for the airline in terms of missed revenue opportunities. Joint optimization of price levels taking into account the diversion percentages believed to exist in the market can suggest the preferred pricing strategy for a carrier to pursue under a given fare product structure.

The gap between the two total revenue curves shown in Figure 5.24 presents the difference in revenues between the airline admitting and denying the existence of passenger diversion behavior in the base case of the constant cost model with linear demands, varying d_{12} (the diversion rate between fare products 1 and 2). The revenue difference is quite small for low diversion percentages. As diversion increases, however, the revenue gap increases. In the limit, over a 14% improvement in revenue exists when the carrier admits that 100% diversion occurs from fare product 1 to 2 in the example.

The price level optimization for the base case varying d_{12} increases the optimal price levels for fare products 1, 2, and 3 shown in Figure 5.11 when the proper level of passenger diversion is predicted. The percentage difference between the fare product 1 optimal price level in the two cases is 15.4% for a 95% diversion rate between fare products 1 and 2 compared to the no diversion case as shown in Table 5.2. Percentage differences of 58.7% and 66.6% exist between the optimal price levels of fare product 2 and 3, respectively, for a 95% fare product 1 to 2 diversion rate relative to the no diversion case. When passenger diversion is ignored, the price levels are held to values lower than is optimal for the airline (*i.e.*, at the levels in the zero diversion rate case in Figure 5.11). The result of ignoring passenger diversion for the airline in the first diversion case is a loss of revenue caused by pricing the fare products too low (*e.g.* below the optimal price levels for the amount of diversion occurring in the market).

Airline Policy	Ignore Diversion	Admit Diversion	Percent Difference
Percentage Diversion	95%	95%	
Revenue	$25691	$29519	+14.9%
FP1 Price	$406	$469	+15.4%
FP2 Price	$213	$337	+58.7%
FP3 Price	$98	$156	+66.6%
Consumer Surplus	$17357	$9813	-43.5%
Passengers	135	115	-15.4%

Table 5.2 Effects of Ignoring and Admitting Diversion on the Base Case with 95% Diversion between Fare Product 1 and Fare Product 2

Base Case Varying d12 - Total Revenue : Admitting vs. Ignoring Diversion
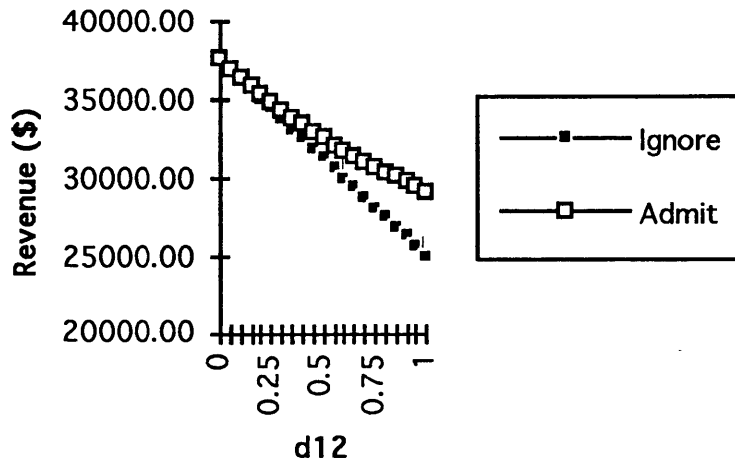
Figure 5.24 - Base Case Varying d_{12} - Total Revenue: Admitting vs. Ignoring Diversion

The joint price level optimization performed under the assumption that passenger diversion exists allows the airline to improve its revenues by charging prices that reduce the revenue dilution caused by diverting passengers. The revenue improvement is achieved, however, at the expense of consumers as witnessed by the decrease in consumer surplus shown in Figure 5.25. In effect, the reoptimization of the joint price level problem shifts much of the gain in consumer surplus resulting from diversion back to the airline. Total market welfare is maintained at a constant level as long as the total number of seats sold does not change leading to a one-to-one tradeoff between airline revenue and consumer surplus.

5.3.5 Effects of Passenger Diversion on Introducing More Fare Products

This subsection discusses the measurement of the effects of offering more fare products on market revenues and consumer surplus when passengers can divert to more-restricted fare products. Overall, the existence of passenger diversion in the market only serves to lower the total revenues earned by the carrier resulting in a lesser revenue gain for additional fare products than when no diversion may occur *ceteris paribus*. The effect of introducing more fare products when passenger diversion exists in the market, however, is not determined. The interactions between the fare products determine whether the introduction of an additional fare product will have a positive or negative revenue effect. For instance, a fare product that has a lower price than any one currently offered with virtually no attached purchase restrictions would almost certainly have a negative revenue impact on the market resulting from diversion increases.

The introduction of an additional fare product to the market may result in positive incremental revenues if the passenger diversion percentages are small and do not change markedly resulting from the introduction of the new fare product. In the limit, the optimal solution with infinitesimally small passenger diversion percentages approaches the no diversion result. In this case, the introduction of an additional fare product is likely to be revenue positive. On the other hand, if many more passengers are able to divert as a result of the ineffective fence that accompanies the newly introduced fare product, it is possible that introducing an additional fare product will have a negative effect on revenue. In either case, measuring the effects of adding fare products to an

Base Case Varying d12 - Total Consumer Surplus : Admitting vs. Ignoring Diversion
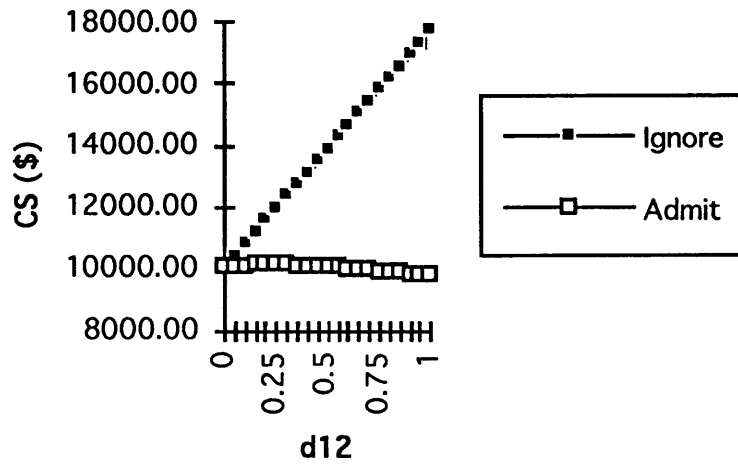
Figure 5.25 - Base Case Varying d$_{12}$ - Consumer Surplus : Admitting vs. Ignoring Diversion

OD market must be made using predictions about the nature of passenger diversion that occurs (or will occur) between the fare products.

The introduction of a new fare product to a market in which passenger diversion is present may result in increased or decreased revenues in contrast to the guaranteed positive revenue impact occurring in the no diversion case. Thus, it is important to incorporate a proper measure of passenger diversion into the generalized cost model since its incorrect omission can change the preferred strategy of an airline.

5.4 Linear Constant Cost Model with Diversion : Applying Booking Limits

In this section, booking limits are applied to the constant cost model of airline fare product differentiation to demonstrate the effects on the optimal solution to the joint price level optimization of capacity controls on the individual fare products. Booking limits are applied to the constant cost model with diversion and used to provide a look at passenger sell up behavior in the presence of capacity constraints.

Investigating the effects of capacity controls on demand, revenue, and welfare can help to provide insight into the methods used by airlines to mitigate the revenue dilution problem. In this section, the application of booking limits to the constant cost model with diversion provides a more realistic look at the passenger demand situation facing commercial airlines. While the analysis performed here is limited to the static case, there is benefit to exploring the effects on demand and revenue of the application of booking limits to the generalized cost model with diversion. Insights concerning passenger behavior in the presence of booking limits are of particular value to airlines when considering changes in the OD market fare product structure.

5.4.1 Measurement of Passenger Diversion Recovery Potential

The revenue losses incurred by the airline when passenger diversion exists in the market can be substantial and airlines seek to minimize them. The previous section demonstrated one method to reduce diversion losses : by admitting that

240

diversion is a problem and performing a reoptimization of prices based upon the level of passengers diversion expected to occur in the market. Applying booking limits to the fare products provides another method for airlines to curtail losses resulting from passenger diversion.

Reversing the negative revenue impact of passenger diversion using booking limits can be achieved by inducing the diverting passengers to purchase higher-priced fare products or, ideally, the fare products for which the airline has targeted them. In other words, booking limits are designed to reverse passenger buy down behavior. For each individual fare product i, the number of diverting passengers from that fare product who may be induced to sell up from lower-priced fare products j by applying booking limits is :

$$\text{Diversion Recovery Potential}_i = \sum_{j=i+1}^{N} d_{ij}Q_i \tag{5.35}$$

For the OD market, the number of passengers potentially selling up is simply equal to the total number of diverting passengers or :

$$\text{Total Diversion Recovery Potential} = \sum_{i=1}^{N-1} \sum_{j=i+1}^{N} d_{ij}Q_i \tag{5.36}$$

since passengers targeted to the lowest-priced fare product (N) cannot divert to a lower-priced fare product.

Each diverting passenger results in lost revenue for the airline. The revenue loss to the airline as a result of diversion for each fare product i is :

$$\text{Revenue Loss}_i = \sum_{j=i+1}^{N} (P_i - P_j)d_{ij}Q_i \tag{5.37}$$

since the airline was originally expecting to earn P_i for each diverting passenger targeted for fare product i but instead received only P_j. The resulting revenue loss for the OD market is :

241

$$\text{Total Revenue Loss} = \sum_{i=1}^{N-1} \sum_{j=i+1}^{N} (P_i - P_j)d_{ij}Q_i \tag{5.38}$$

It is clear from equation 5.38 that the revenue loss is amplified by increased diversion percentages and large price differentials. The total revenue lost by diversion is the maximum amount that can be recovered by applying booking limits to the solution.

Similarly, the deadweight welfare loss associated with passengers diverting from each fare product i to the lower-priced fare products is :

$$\text{Welfare Loss}_i = \sum_{j=i+1}^{N} \sum_{i<k \leq j} c_k d_{ij}Q_i \tag{5.39}$$

making the total loss for the OD market :

$$\text{Total Welfare Loss} = \sum_{i=1}^{N-1} \sum_{j=i+1}^{N} \sum_{i<k \leq j} c_k d_{ij}Q_i \tag{5.40}$$

Equation 5.40 demonstrates that the welfare loss results exclusively from the degradation costs associated with the lower-priced fare products incurred by the diverting passengers. Larger numbers of diverting passengers and high values of degradation costs result in increased welfare loss and thus, recovery potential. In fact, with zero degradation costs for the lower priced fare products, the welfare loss is eliminated. Again, society would benefit from a reduction in degradation costs.

Passenger diversion results in a consumer surplus gain at the expense of airline revenue and societal welfare. For any fare product i, the consumer surplus gain associated with passenger diversion is :

$$\text{Consumer Surplus Gain}_i = \sum_{j=i+1}^{N} \left[P_i - \sum_{i<k \leq j} c_k - P_j \right] d_{ij}Q_i \tag{5.41}$$

242

The measure of changing consumer surplus simply represents the shift of revenues away from the airline to the consumer minus the deadweight welfare loss to the market associated with the increased degradation costs. An increase in degradation costs results in a decreased consumer surplus gain. Not only do increased degradation costs reduce optimal welfare and revenue increasingly but also decrease the amount of consumer surplus. Thus, degradation costs are a detriment to all *ceteris paribus*.

The total gain of consumer surplus for all fare products offered in the origin to destination market is :

$$\text{Total Consumer Surplus Gain} = \sum_{i=1}^{N-1} \sum_{j=i+1}^{N} \left[\left(P_i - \sum_{i<k \leq j} c_k - P_j \right) d_{ij} Q_i \right] \qquad (5.42)$$

Consequently, the maximum consumer surplus that can be recovered by the airline is shown in equation 5.42. It is interesting to note that with zero degradation costs for the lower-priced fare products, the consumer surplus gain and the revenue loss are equal (*e.g.* no welfare is lost).

5.4.2 Passenger Switching Behavior in the Presence of Booking Limits

Revenue recovery from the application of booking limits can only result from the inducement of diverting passengers to purchase higher-priced fare products or "sell up". In the context of booking limits, an increase in the revenue earned by the airline through passenger sell up behavior may only result from the denial of lower-priced fare products to the diverting passengers. To begin, passenger fare product switching behavior must be defined to properly measure the effect of booking limits on the solution to the joint price level optimization. The measurement of revenue, welfare, and consumer surplus impacts can then be explored under the assumed switching behavior.

All diverting passengers are willing to purchase the fare products for which they were initially targeted by the airline. When allowed to divert, however, they prefer to purchase the lowest-priced fare products for which they are able to meet the restrictions. When booking limits force diverting passengers to sell up, the general rule for passengers in terms of sell up behavior is that they purchase

the lowest-priced fare product available to them under the prevailing booking limits. Diverting passengers denied a seat at the price level of the lowest-priced fare product that they are able to purchase seek to purchase the next-lowest-priced fare product until availability enables them to book a fare product. Consumers attempt to purchase fare products up to and including the fare product for which they were targeted by the airline. If the consumer is denied his targeted fare product and all lower-priced fare products that he is eligible to purchase, he will choose not to travel.

A simple example using a three fare product example with diversion from fare product 1 to fare product 3 (exclusively) highlights the sell up and diversion behavior undertaken by consumers in the presence of booking limits. With no binding booking limits, passengers diverting from fare product 1 will purchase fare product 3. Diverting passengers unable to purchase fare product 3 because of an exceeded booking limit will purchase fare product 2 if available. Recall that passengers able to meet the restrictions associated with fare product 3 are guaranteed to be able to meet the fare product 2 restrictions and have been assumed to prefer the lowest priced fare product for which they can meet the restrictions. If, however, fare product 2 is unavailable due to booking limits, the diverting passengers will purchase fare product 1 (for which they were targeted by the airline). It is never logical for the airline to place booking limits on the unrestricted fare product in the generalized cost model since passengers cannot divert to it. Thus, fare product 1 provides no sell up opportunity.

Airlines using booking limits also risk denying passengers service who normally would be willing to travel at the prevailing price levels but are not willing to sell up to the higher-priced fare products. For example, a passenger hoping to divert from fare product 2 to fare product 3 who is denied a seat for both fare products 3 and 2, does not have a willingness to pay high enough to elicit the purchase of fare product 1. Thus, this fare product 2 passenger represents a lost passenger booking instead of a candidate for sell up when fare product 2 and 3 booking limits are binding. The potential for service denials introduced by booking limits makes their revenue impacts uncertain.

5.4.3 Recovery of Diversion Losses Using Booking Limits

The order of passenger bookings within the single booking period assumed in
the constant cost model determines the amount of sell up behavior versus
passenger service denials occurring in the market when booking limits are placed
on the individual fare products. Calculations of shifts in the market performance
measures of revenue, welfare, and consumer surplus, in the presence of booking
limits are also affected by the order of passenger bookings and the assumed
behavior of passengers. This subsection addresses the recovery potential for
these market performance measures when booking limits are applied.

Booking limits which prevent the diversion of all passengers without causing
passengers not to travel would allow the airline to capture all of the sell up
potential in the market without diluting revenue. In theory, a complete recovery
of revenue (and welfare) potential through the application of booking limits is
possible. This can only be achieved, however, when the price levels prevailing in
the market are the optimal price levels for the no diversion case. This is logical
since the induced sell up of all diverting passengers would lead to a replication
of the optimal solution to the no diversion case.

The sole purpose of applying booking limits to the generalized cost model is to
prevent passenger diversion. Accordingly, the revenue recovery provided by the
application of booking limits results from the prevention of passenger diversion.
Unfortunately, there is a danger that the use of booking limits to prevent
passenger diversion may lead to denied passenger bookings. Airlines face a
tradeoff between the possibility of denying passengers service and inducing
diverting passengers to sell up when booking limits are applied to the
generalized cost model. A positive revenue contribution from the application of
booking limits to the generalized cost model may only be achieved if the
incremental revenues gained from passengers denied seats for lower-priced fare
products through induced sell up exceeds the revenue loss from passengers
targeted to the lower-priced fare product who are denied seats and consequently
decide to forego travel.

While the losses associated with diversion are independent of the arrival process
in the market, the recovery of these losses through the application of booking

limits is not. In other words, the effects of passenger sell up behavior on revenue depends upon the arrival discipline in the market. The amount of revenue and welfare recovered may vary widely for the same market conditions with different arrival processes. The lost revenue and welfare can only be gained back by the airline and society when a correct application of fare product booking limits is used with an amenable passenger arrival discipline within the single booking period. The passenger arrival process may circumvent the revenue recovery tactics attempted by the airline.

Limited bounding analysis can be performed when an arrival process is assumed within the single static period. The passenger denial/sell up tradeoff can be bounded by the cases of passenger bookings in increasing and decreasing order of willingness to pay. Passenger arrivals in increasing order of willingness to pay can achieve optimal system revenues equal to the optimal no diversion case while passengers arrivals in decreasing order cannot. When passengers arrive in increasing order of willingness to pay, a binding booking limit associated with a single fare product induces all diverting passengers to sell up while not denying service to any passengers targeted to the restricted fare product. Passenger arrivals in decreasing order of willingness to pay, on the other hand, result in the denial of service to passengers targeted to the lower-priced fare product without inducing a single diverting passenger to sell up. Consequently, bookings in increasing order of willingness to pay are more effective in helping the airline to achieve its revenue maximizing goal when booking limits are used in this way.

5.4.4 Passenger Bookings in Increasing Order of Willingness to Pay

The optimal solution to the no diversion case (resulting in maximum system revenue) can be replicated in the case with diversion when booking limits are applied and passengers arrive in increasing order of willingness to pay. When the booking limits are set equal to the total number of passengers for each fare product in the optimal no diversion case, the optimal solution to the case with diversion replicates the no diversion optimal solution. All revenues, consumer surplus, and societal welfare are restored to their levels in the no diversion case. In other words, the case of passenger arrivals within the single booking period in increasing order of willingness to pay can result in an optimal solution equal to that in the no diversion case when optimal booking limits are applied.

Specifically, the decreased welfare and revenue resulting from the diverting passengers are returned to their no diversion levels while the consumer surplus gains made by the diverting passengers are removed.

Figure 5.26 presents the base case varying d_{12} in which the entire 50% of diverting fare product 1 passengers sell up to the higher-priced unrestricted fare product as a result of booking limits being placed on fare product 2. The hatched area represents the deadweight welfare loss recovered due to the diverting passengers being forced to sell up. The lightly shaded area shown in the figure represents the consumer surplus that is lost when the passengers diverting to fare product 2 are forced to sell up to fare product 1. Finally, the combination of the consumer surplus and recovered welfare areas represents the revenue returned to the airline.

As demonstrated in the previous section, higher revenues are earned by the airline when diversion is admitted by the airline and optimal price levels are set assuming a diversion percentage occurring in the market. Table 5.3 shows the greater optimal revenues ($32573) for the 50% diversion base case in which the airline admits diversion compared to when diversion is ignored ($31322). When optimal booking limits are applied to the model, however, the opposite is true.

Airline Policy	Admit Diversion	Ignore Diversion	Admit Diversion	Ignore Diversion
Booking Limits	Perfect	Perfect	None	None
Revenue	$36577	$37575	$32573	$31322
FP1 Price	$406	$445	$406	$445
FP2 Price	$213	$290	$213	$290
FP3 Price	$94	$133	$94	$133

Table 5.3 Comparison of Perfect Booking Limit Applications : Passengers Book in Increasing Order of Willingness to Pay

The optimal solution to the base case with d_{12} diversion in which the airline reoptimizes price levels (admit diversion) <u>and</u> optimal booking limits are applied (which eliminate diversion in the market altogether) provides a total revenue of
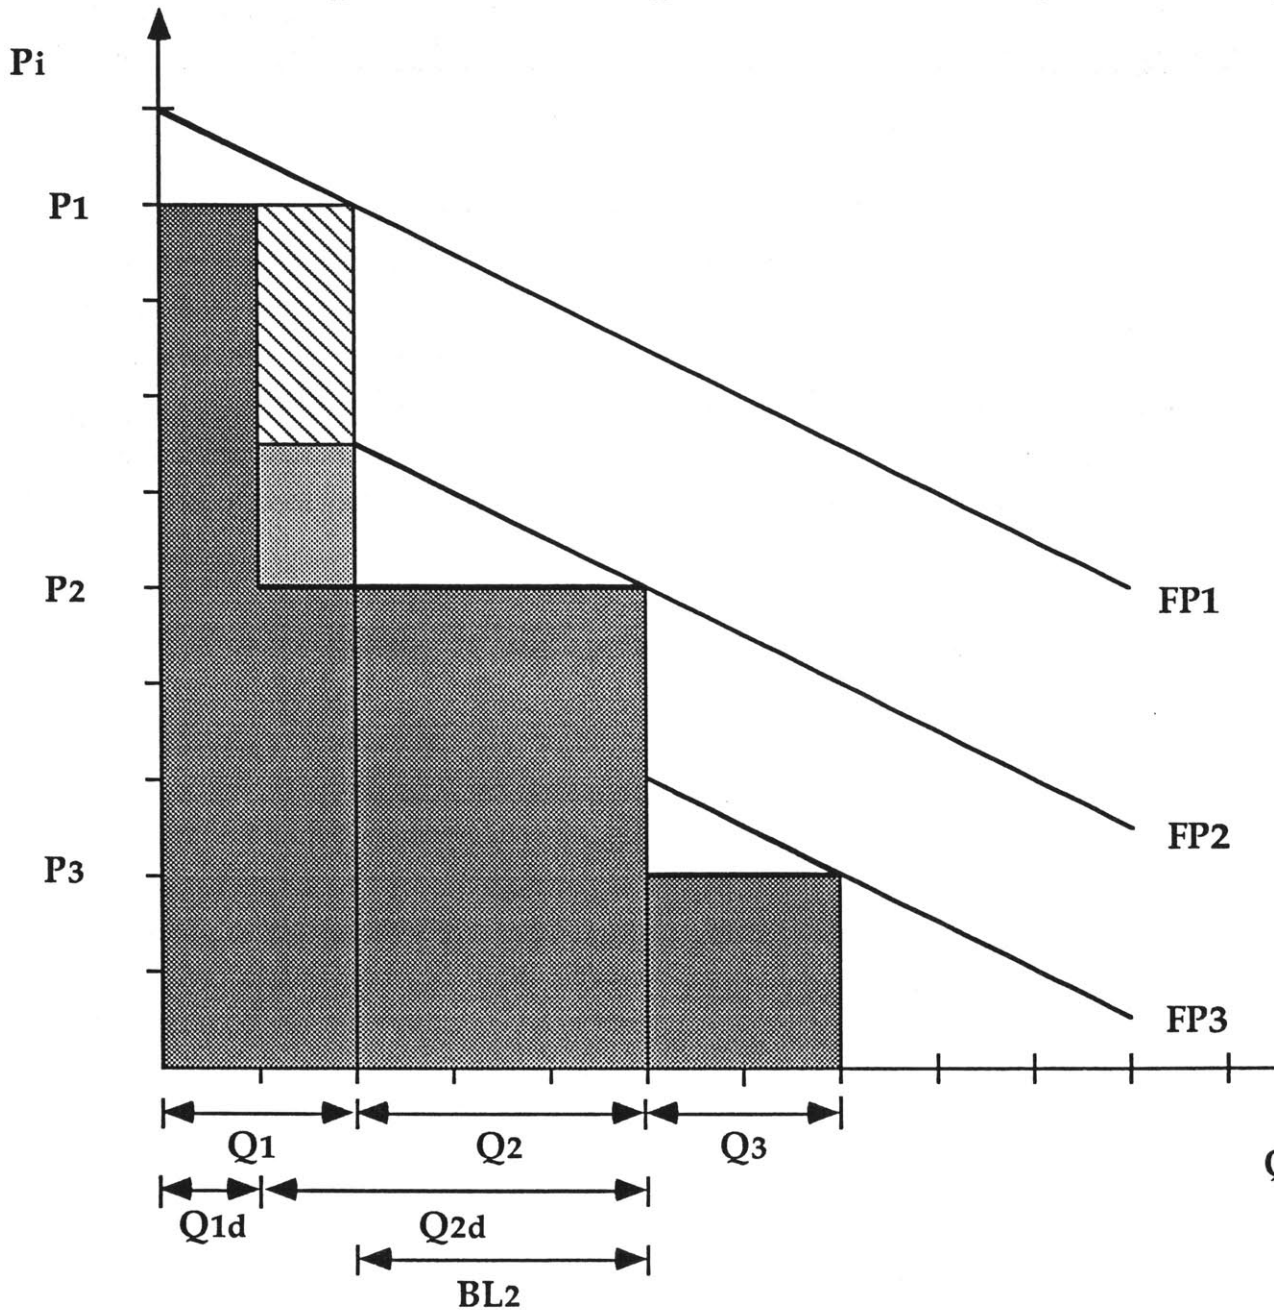
Figure 5.26 - Generalized Cost Model With d₁₂ Diversion at 50% - Three Fare Product Example : Booking Limits Application With Bookings in Increasing Order of Willingness to Pay

$36577 as shown in Table 5.3. In the case where d_{12} diversion is not admitted (ignore diversion) and booking limits are applied that eliminate diversion, airline revenues are equal to $37575. The case in which the airline has both reoptimized price levels and applied optimal booking limits to correct diversion has resulted in lower revenues than when optimal booking limits are applied without a reoptimization of price levels. This implies that price level reoptimization admitting diversion, while providing increased revenues in the absence of booking limits, is likely to result in a lower revenue potential for the market.

Therefore, the airline faces a choice when the techniques of booking limits and acknowledgment of diversion are available to it. The airline may attempt to optimize revenues accepting that passenger diversion is going to proliferate or it may attempt to curb diversion by applying booking limits. A combination of the two methods may also be chosen as the preferred strategy, however, any admission of diversion may render the optimal solution unattainable. Blindly applying both diversion-correction methods, however, may lead to lower revenues.

5.4.5 Passenger Bookings in Decreasing Order of Willingness to Pay

For passenger arrivals in decreasing order of willingness to pay, on the other hand, the optimal solution to the no diversion case can never be replicated when diversion exists regardless of the booking limits applied. When booking limits are applied, the diverting passengers are always the first to arrive for booking. The diverting passengers with the highest values of willingness to pay are, thus, guaranteed first chance at securing seats for the lower-priced fare products. The passengers targeted for the lower-priced fare products book after all diverting passengers have already made their purchases. Recall that the targeted passengers cannot be induced to sell up to higher-priced fare products. Therefore, with passenger bookings in decreasing order of willingness to pay, no revenue can be recovered by forcing diverting passengers to sell up without rejecting all passengers targeted to the lower-priced fare product.

Figure 5.27 shows the revenue and welfare lost by the application of booking limits (the optimal booking limits for the case in which passengers arrive in increasing order of willingness to pay) when passengers arrive in decreasing

249

Generalized Cost Model With d12 Diversion at 50%
Three Fare Product Example
Booking Limits Application With Bookings
in Decreasing Order of Willingness to Pay
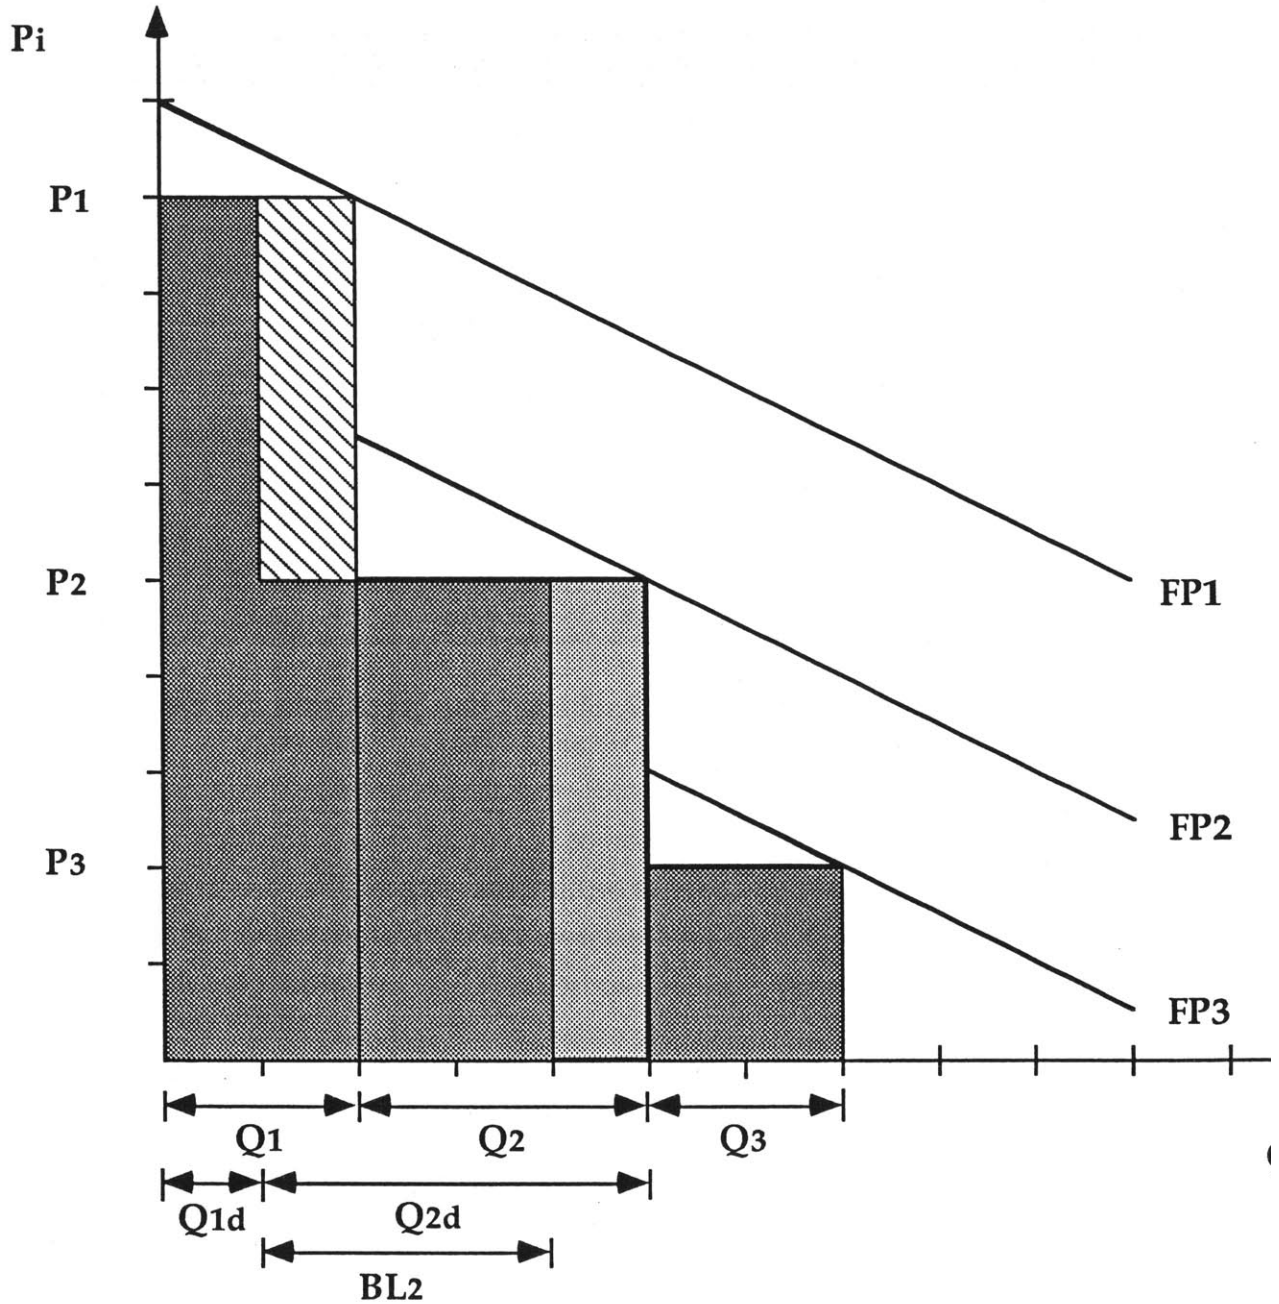
Figure 5.27 - Generalized Cost Model With d$_{12}$ Diversion at 50% - Three Fare
Product Example : Booking Limits Application With Bookings in Decreasing
Order of Willingness to Pay

order of willingness to pay. The lightly shaded area in the figure represents the revenue loss associated with the denial of service to the passengers initially targeted to fare product 2. The revenue dilution associated with diversion from fare product 1 to 2 (shown in the hatched area of Figure 5.27) still plagues the airline and the booking limits do nothing to prevent it. In fact, the booking limits only serve to worsen the revenue performance of the carrier by denying seats to passengers targeted to fare product 2. Total revenue is equal to $25081 in the case with booking limits in contrast to the $32573 total revenues for the case without booking limits.

Airline Policy	Admit Diversion	Ignore Diversion	Admit Diversion	Ignore Diversion
Booking Limits	Yes : As in Table 5.3	Yes : As in Table 5.3	No	No
Revenue	$25081	$24460	$32573	$31322
FP1 Price	$406	$445	$406	$445
FP2 Price	$213	$290	$213	$290
FP3 Price	$94	$133	$94	$133

Table 5.4 Comparison of Perfect Booking Limit Applications : Passengers Book in Decreasing Order of Willingness to Pay

With the counterproductive application of booking limits performed when bookings occur in decreasing order of willingness to pay, the case in which price levels are reoptimized admitting the existence of diversion outperforms the case in which diversion is ignored (as shown in Table 5.4) in contrast to when passenger arrivals occur in increasing order of willingness to pay. When booking limits are applied to the base case varying d_{12}, the actual revenue total is expected to fall somewhere below the upper bound of $37575 (e.g. the no diversion optimal solution) but the exact result depends upon the specific arrival process and booking limits actually occurring in the market.

5.4.6 Booking Limits Summary

In summary, the impact of booking limits on the optimal solution to the joint price level optimization is not determined unless the booking arrival process of passengers within the single static booking period is known. The revenue result can, however, be bounded from above assuming, as before, an arrival process in increasing order of willingness to pay. The lower bound on revenue can drop as low as zero with an illogical application of booking limits, however, a more realistic lower bound can be set depending upon the existing market conditions. The total revenue recoverable (sell up potential) from the application of booking limits remains limited to the amount of revenue lost from passenger diversion. With an unknown booking arrival process, the benefits of reoptimizing price levels based on the amount of diversion assumed to exist in the market are unclear as well when booking limits are introduced. Thus, knowledge of the booking process has a direct impact on the chosen airline price setting and booking limit strategy.

5.5 Applications of the Generalized Cost Model

The generalized cost model can be used for other airline marketing applications in addition to the analysis of the relationships between price levels, restrictions, and other market conditions. The section begins by using the model to evaluate the revenue impacts of different aircraft configurations on the flight in the market. The impact of introducing fare products into the existing market fare product structure is then explored. The constant cost model is once again the chosen specification but, this time, constant elasticity demand functions are also used for some of the analysis. As before, other model formulations and demand specifications may be used in virtually the same manner with slight modifications.

5.5.1 Aircraft Configuration Application

The generalized cost model can be used to measure the effects of a variation in a cabin configuration when more than one class of service is offered on the aircraft. Since a primary difference in cabin service level is the size of the seat, increasing

252

the number of premium class seats decreases the total capacity of the aircraft. In this way, the cabin configuration can directly affect revenue. The lowered capacity resulting from an increased number of premium class seats must be traded off with the revenue benefits of the price levels paid by first and business class passengers. The total revenue impacts determined by applying the model can be used to evaluate the preferred cabin configuration when more than one configuration or aircraft type is available (or proposed).

A three class cabin configuration example is used to demonstrate the uses of the constant cost model as a tool for the evaluation of airline marketing planning decisions. Unrestricted fare products are offered to consumers in each of the three different cabin classes of service :

1) First Class (F)
2) Business Class (J)
3) Coach Class (Y)

Three additional constraints must be added to the formulation of the joint price level optimization problem to account for the fixed cabin capacities :

$$Q_F \leq Q_{Fmax} \tag{5.43}$$

$$Q_J \leq Q_{Jmax} \tag{5.44}$$

$$Q_Y \leq Q_{Ymax} \tag{5.45}$$

where $Q_{Imax} =$ the capacity of cabin I

The constant cost model is evaluated assuming a constant elasticity demand function in the three cabin class example. The general functional form for demand in the constant elasticity case appears below as input to the constant cost model :

$$P_i = a \left(\sum_{k=1}^{i} Q_k \right)^b - \sum_{k=1}^{i} c_k \tag{5.46}$$

253

The formulation of the generalized cost model is identical to that presented in section 5.1 of the dissertation with the addition of constraints 5.43, 5.44, and 5.45. The three cabin base case has the following OD market parameter values :

1) a = 2500
2) b = -3
3) Cap = 293
4) Q_{Fmax} = 35
5) Q_{Jmax} = 134
6) Q_{Ymax} = 124
7) c_F = 0
8) c_J = 50
9) c_Y = 50

Thus, the demand functions for the three fare products in the three different cabins for the OD market are :

$$P_F = a(Q_F)^b \tag{5.47}$$

$$P_J = a(Q_F + Q_J)^b - c_J \tag{5.48}$$

$$P_Y = a(Q_F + Q_J + Q_Y)^b - (c_J + c_Y) \tag{5.49}$$

in the three cabin class base case for fare products F, J, and Y, respectively.

The degradation costs associated with each fare product represent the decrease in level of service from first to business and finally, to coach class. It may be more intuitive to view the degradation costs as the improvement in level of service from coach class to business class to first class. In the example, the business and first class services are valued at $50 and $100 more than coach class, respectively, by all passengers.

Neither booking limits nor passenger diversion are assumed to exist in the three cabin class base case. Mathematically, the following parameters apply in the constant elasticity three cabin class base case formulation :

$$
\begin{aligned}
&1)\ d_{FJ} = 0\\
&2)\ d_{FY} = 0\\
&3)\ d_{JY} = 0\\
&4)\ BL_F = \infty\\
&5)\ BL_J = \infty\\
&6)\ BL_Y = \infty
\end{aligned}
$$

The optimal values of price level with resulting passenger demand and airline revenue by fare product (and cabin) appear below in Table 5.5 for the three cabin classes :

Product	Passengers	Price	Revenue
F	35	$764	$26753
J	134	$402	$53902
Y	124	$276	$34282

Table 5.5 Three Cabin Class Base Case

The two aircraft configuration applications tested using the constant cost model with constant elasticity demands are the measurement of the effects of service improvements and the variation of actual cabin configurations on revenue. The model can be used by airlines to measure the expected revenue impacts of their proposed marketing decisions providing decision support to marketing planners. The effects of service improvements on revenue are explored first followed by the impacts of variations in cabin configuration.

Measuring the Effects of Service Improvements on a Fixed Cabin Configuration

In the short term, the configuration of an aircraft must remain fixed. Obviously, seats cannot be interchanged on a flight by flight basis. Therefore, it is of interest to the airlines to know the effects of the fare product structure on revenues for a fixed aircraft cabin configuration. In this vein, airlines considering introducing service improvements in OD markets have interest in knowing the revenue benefits or hazards of their actions. Service improvements are defined as increases in the amenities provided to a cabin class such as upholstering seats

255

with leather, installing cellular phones in seat backs, or garnishing each salad with an additional olive.

Prior to making a fixed cost service improvement to an aircraft cabin, an airline would like to determine the number of trips required to recuperate the capital outlay. The ability to judge the length of time required to pay off a service investment would aid marketing planners in their assessment of fixed cost investment decisions. An airline also has interest in knowing whether or not the revenue return justifies a service improvement requiring a recurring investment. Airline marketers would value a measurement of the recurrent benefit or loss expected from a proposed service improvement when making resource allocation decisions. The generalized cost model quantifies the revenue impacts of both fixed and recurrent cost service improvements for different cabin classes under the optimal (or any other) set of price levels.

To perform a cost-benefit evaluation of service improvements in a particular OD market, the effect on the optimal fare product price levels of making a fixed or recurrent cost investment is also of interest to airline planners. For instance, if an airline is unable to charge price levels that approximate the optimal, the length of time required to recover the fixed cost investment increases. The expected revenue returns to service improvements under a given set of price levels provided by the generalized cost model allows airline planners the ability to measure the expected revenue impacts of service improvements taking into consideration the fare product structure and pricing decisions being made in a particular market. The revenue effect information can be used to evaluate the advisability of investing in proposed service improvements.

In the three cabin class base case example, consider first a fixed cost investment of $1,500,000 for the improvement of the business class cabin. For instance, the airline may plan to install a personal video screen in each seat back along with a fax machine for the exclusive use of the business class cabin. The proposed service improvements are expected to result in a decrease in the degradation costs associated with the business class cabin by $25 for all passengers relative to first class. The constant cost model can be used to determine the impact of the change on the OD market.

The results of the service improvements on the optimal price levels with resulting demand and revenue are shown in Table 5.6. At optimality, the total flight revenue for the base case of the three cabin configuration example equals $114,938. When the service improvement is made, optimal flight revenue is improved by $3350 per flight to $118,288. Thus, at a minimum, it would take the airline 450 flights to earn back the fixed cost required for the service level improvement.

Case	Base	More Service
C2	$50	$25
Revenue	$114938	$118288
F Price	$764	$764
J Price	$402	$427
Y Price	$276	$276
F Passengers	35	35
J Passengers	134	134
Y Passengers	124	124

Table 5.6 Fixed Cost Service Improvement

The optimal price level for the business class cabin increases as a result of the decreased degradation costs from the fixed cost improvement. The optimal price levels for the first and coach class cabin, on the other hand, remain the same after the service improvement. Since the example is capacity constrained, the number of passengers carried at optimality does not shift with the decreased degradation costs. Thus, the degradation cost reduction translates directly into increased airline revenue from the increased business class price levels.

A variable cost investment may also be investigated using the generalized cost modeling framework. The airline marketing department considers a service improvement which requires a recurrent investment of $20 per passenger per flight in the business class cabin. For instance, passengers may be accommodated in the airline's members only or first class lounge prior to departure. The degradation costs associated with the business class fare product are expected to be reduced by $10 when the service improvement is made. The

257

airline has chosen to make the improvement if the recurrent investment has a positive net impact on system revenues.

The impact of the service improvement on optimal price levels and the resulting revenue and demand levels are presented in Table 5.7. Total market revenue would be improved by $1340 dollars (from $114,938 to $116,278) when the service improvement is put in place. The total number of business class passengers carried is equal to 134 in the example making the total expected cost of the service improvement per flight $2680. Since the revenue cost outweighs the revenue benefit, a negative net revenue contribution would be made by the service improvement and thus, the change should not be made.

Case	Base	More Service
C2	$50	$40
Revenue	$114938	$116278
F Price	$764	$764
J Price	$402	$412
Y Price	$276	$276
F Passengers	35	35
J Passengers	134	134
Y Passengers	124	124

Table 5.7 Variable Cost Service Improvement

It should be noted that the cost of the variable cost service improvement is simply twice the benefit and this is translated directly into the optimal solution. The simplicity of the results comes from the binding capacity constraint in the optimal solution. A decrease in degradation costs for a single fare product when the capacity constraint is binding generally results in a decrease in the optimal price level for that fare product only. The generalized cost model would provide less trivial results for capacity unconstrained cases. The limitations of the constant elasticity formulation, however, required that the capacity constrained case be used to present realistic results. Thus, in the interests of demonstrating the flexibility of the generalized cost model, somewhat simplistic test case results are presented.

Measuring the Effects of Varying the Cabin Configuration

It is also of interest to airlines to evaluate different cabin configurations for a given route if, for instance, there is flexibility in the assignment of aircraft to the market. Clearly, the analysis used to evaluate the cabin configuration of an aircraft serving multiple routes with multiple demands becomes increasingly complex. Nonetheless, the evaluation of cabin configurations using the generalized cost model can lead to improved aircraft assignments for a given market when there is scheduling flexibility of a common aircraft type with different configurations. It should be noted that similar analysis can be performed when airlines are able to switch aircraft types.

In the three cabin example, the aircraft cabin configuration is varied to demonstrate the effects on revenue of trading off seats between the coach and premium class cabins. The Boeing 747-200 series aircraft operated by United Airlines are configured in two different ways and provide a realistic example of cabin configuration differences within the same airline. The 747-200 YI configuration has 35 first class seats, 134 Connoisseur® (business) class seats, and 124 economy class seats (OAG, 1993). Coincidentally, the United 747-200 YI configuration matches the three cabin class example base case exactly. The 747-200 YR configuration, on the other hand, has 18, 79, and 266 seats in the three cabin classes, respectively. The YI configuration has a total of 293 seats while the YR has 363 total seats.

The use of the YI configuration increases the per seat revenue potential (with higher numbers of first and business class seats) but decreases the total number of seats available on the aircraft. Consequently, the revenue effects of using each configuration are dependent upon the market conditions. The generalized cost model can be used to determine the preferred aircraft configuration in terms of revenue contribution on a market by market basis. The results derived from the optimal price levels for the two United 747-200 aircraft configurations in the three cabin class base case varying configuration type appear in Table 5.8.

Aircraft	747-200 YI	747-200 YR
Revenue	$114938	$122854
F Price	$764	$954
J Price	$402	$494
Y Price	$276	$251
F Passengers	35	18
J Passengers	134	79
Y Passengers	124	266
Capacity	293	363

Table 5.8 Comparison of United Airlines 747-200 Aircraft Configurations

The optimal revenues available for the YI configuration in the three cabin base case are $114,938 while they are $122,854 for the YR configuration. The YR configuration dominates the YI configuration in terms of optimal revenue performance for the specific market analyzed. The optimal revenue returns only result, however, when the carrier charges optimal price levels in the market.

The optimal first and business class price levels for the YR configuration are $190 and $92 dollars greater than for the YI, respectively. Thus, the optimal price levels are much higher for the premium class passengers when the YR configuration is used. The relative scarcity of premium class seats in the YR configuration leads them to be more valuable to consumers with the high level of demand that exists in the example market. The optimal coach class fare level, on the other hand, is $25 lower for YR configuration. This results primarily from the high level of market demand (all results are capacity constrained) and the greater number of passengers that may be accommodated aboard the YR aircraft.

From purely an optimal revenue perspective, the YR configuration should be used to serve the market when having aircraft assignment flexibility and dictating market price levels. Inability to set, and have passengers pay, price levels approximating the optimal ones, however, may change the preferred aircraft configuration for the market. In either case, the generalized cost model may be used to evaluate the preferred aircraft type or configuration for any set of market price levels.

The important interrelationship between capacity provision and market price levels can be addressed in the generalized cost modeling framework. The model provides a method to incorporate the effect of pricing decisions and demand characteristics into capacity allocation decisions. The incorporation of the impact of market price levels into aircraft capacity decisions provides the foundation for the development of planning tools that synthesize two of the most important areas of marketing planning.

The examples presented here presume a fixed aircraft configuration which can be evaluated from a revenue impact perspective. Optimizing the actual configuration for a given type of aircraft using the generalized cost modeling framework requires a reformulation of the problem that introduces cabin sizes as decision variables. Additional constraints would need to be added to ensure that the optimal cabin configurations are spatially feasible for the fixed aircraft size. Although such an application is of great interest to airline planners, it is left to future research efforts.

5.5.2 Proposed Fare Product Introductions and Modifications

The marketing departments of airlines attempt to segment the market and stimulate demand through the introduction of different types of fare products onto the market. The decision support tools available at most major carriers that quantify the effects of different fare product introductions are quite limited. The effects of introducing different fare products can be quantified using the generalized cost model helping to fill this decision support void. The information provided can be used by airline marketing managers to determine the most favorable changes to the fare product structure among their proposals.

Varying the fare product mix in a market can have a major effect on the demand and revenue in an OD market. The introduction or removal of a fare product, for instance, can determine whether or not many passengers decide to travel. A fare product that is priced lower than any one currently available on the market can stimulate a large amount of incremental demand. Unfortunately, while such a fare product is intended to stimulate demand, it may only serve to dilute revenues if the segmentation devices (i.e., restrictions) accompanying it are ineffective in the prevention of passenger diversion.

261

The introduction of a new fare product with restrictions similar to an existing fare product can also make an existing fare product redundant on the market. For example, the introduction of a fare product with identical (perceived) restrictions to an existing fare product with a lower price level will drive the demand for the existing fare product to zero. Clearly, no rational consumer would purchase a fare product that is dominated in all attributes by another (ignoring availability constraints). The effect of any fare product introduction is, of course, dependent upon the existing fare product structure and the attributes (restrictions and price level) of the new service offering. The generalized cost model can help to measure the effects of fare product introductions onto the market under any fixed fare product structure.

A simple glance at the generalized cost model formulation allows the identification of where a new fare product fits into the fare product inferiority hierarchy. The degradation costs accompanying the proposed fare product allow for the calculation of a demand function for that fare product from the unrestricted fare product. The degradation cost associated with each prospective fare product must, of course, first be quantified. The position of the isolated demand function in the locus of isolated demand functions can then be determined from the degradation costs relative to the other fare products in the generalized cost model. The model can then be used with the addition of a new fare product or a change to an existing fare product and used to predict the effects on optimal price levels and revenue of the proposed changes to the fare product structure.

The price level of the fare product being introduced to the market can be determined either endogenously or exogenously. When offered pricing flexibility, however, it is preferable to optimally (and endogenously) determine the price level for the fare product to be introduced. Nonetheless, the generalized cost model provides flexibility in the determination of price levels that may prove quite valuable to airline marketing planners.

Comparison of Alternative Fare Products

Imagine that an airline is considering introducing a new fare product onto the market but is unsure which of two potential candidates will have the most

positive (or least negative) impact on revenue. For example, the airline may consider the addition of a three or seven day advance purchase non-excursion fare product to the three fare product base case with linear demands. The degradation cost associated with the three day advance purchase fare product is $25 dollars less than that of the seven day and $25 more than the unrestricted. In terms of the fare product hierarchy, both of the additional fare product candidates fall between the unrestricted and seven day advance purchase excursion fare products.

From the perspective of the airline, introducing the fare product with the lowest degradation costs is always preferable under perfect segmentation. The preferred fare product may differ, however, depending upon the amount of passenger diversion existing in the market. In the example, the expected diversion rate from the unrestricted fare product to the 3 day advanced purchase is assumed to be 40% compared to the 20% assumed for the 7 day advanced purchase. Diversion is not expected to occur between the other fare products. Thus, while passengers prefer the 3 day advanced purchase fare product and demand for it is greater *ceteris paribus*, the increased diversion from the unrestricted fare product may counteract this benefit in terms of revenue performance. The constant cost model base case with linear demand functions is used in the example presented to evaluate which fare product should be introduced by quantifying the total effect on market conditions expected.

Table 5.9 presents a detailed comparison of the two fare product introduction scenarios relative to the existing market fare product structure. Total revenues are greater with the introduction of the 7 day advance purchase fare product which generates $36747 compared to the $36454 earned when the 3 day advance purchase fare product is introduced. The airline would be better served by the introduction of the seven day non-excursion fare product since the negative revenue impacts from the added diversion introduced by the three day advance purchase fare product outweigh the benefits of the lower degradation costs.

263

Fare Product Introduction	None	AP3	AP7
Diversion to New FP	0%	40%	20%
Revenue	$37575	$36454	$36747
Unrestricted FP Price	$406	$471	$453
New FP Price	****	$342	$305
E7NR Price	$213	$186	$179
E14NR Price	$94	$81	$77
Passengers	135	140	141

Table 5.9 Effects of Introducing Non-Excursion Fare Products

As expected, the optimal price level of the AP3 fare product is higher than that of the AP7 when they are introduced separately. The decreased AP3 degradation costs as well as the increased diversion cause the higher optimal price levels for all market fare products relative to the AP7 introduction.

The introduction of either of the fare products, however, leads to a reduction in total revenue from the $37575 earned by the 3 fare product base case without diversion. Thus, the introduction of additional fare products which cause diversion is shown, in this case, to have a negative revenue impact overall. While the seven day advanced purchase fare product performs better than the three day, the introduction of either product would result in lesser revenue performance than currently exists. The potential hazards of introducing fare products that are ineffective at containing passenger diversion has once again been demonstrated using the model.

Restricted Optimization When Price Levels for Certain Fare Products are Fixed

Consider next an airline introducing fare product on the market for which the price level has already been fixed as a result of a system-wide marketing initiative. If a marketing promotion or executive policy by the airline forces the

price level of a certain type of fare product to remain at a specified level, the airline can still optimize the remainder of the fare product price levels since the pricing department still exercises flexibility in their pricing. It should be observed, however, that the optimal airline revenue can never increase (and is likely to decrease) as a result of the additional constraint(s) being added to the joint price level optimization.

The general case entails two subgroups of fare products in which the airline must take as given the exogenously determined price levels of the first subgroup while being able to freely determine the price levels of the second. The airline can optimally set the price levels of the second subgroup given the specified price levels for the first. Price levels can be jointly optimized in a restricted form to yield the optimal price levels of the fare products for which there is pricing flexibility. The methodology is demonstrated below by fixing the price level of the most-restricted fare product in the existing fare product hierarchy of the three fare product constant cost model base case with linear demands.

The restricted optimization problem must be formulated with the added restriction of the exogenously determined fare product 3 price level. In this case, fare product 3 has been increased from the optimal value of $93.75 to $125 as a result of a management attempt to increase fare levels on short haul routes. The constant cost model is used to calculate the optimal price levels for the two fare products (1 and 2) under the direct control of the airline pricing department (*e.g.* not dictated by the fare increase initiative).

The additional constraint added to the formulation of the restricted joint price level optimization appears below for the three fare product base case :

$$P_3 = 600 - 3(Q_1 + Q_2 + Q_3) - 100 = 125 \qquad (5.50)$$

The optimization is performed using the GAMS non-linear optimization package in the same way as before, differing only by the additional constraint on the optimal price level of fare product 3. More complex examples of fare initiatives may place restrictions, for example, on the price differentials of certain products and require more detailed constraint set modifications but can be modeled in the generalized cost framework nonetheless.

The optimal values of price level with resulting passenger demand and airline revenue by fare product appear below in Table 5.10 for the three fare product base case and the related case in which the price level of fare product 3 is constrained to equal $125 :

Case	Base	Constrained
Revenue	$37575	$37361
FP1 Price	$406	$417
FP2 Price	$213	$233
FP3 Price	$94	$125
FP1 Passengers	64.58	61.11
FP2 Passengers	39.58	36.11
FP3 Passengers	31.25	27.78
Passengers	135	125

Table 5.10 Impacts of Constraining the Fare Product 3 Price Level

The fixed $125 price level of the fourteen day advanced purchase fare product leads to optimal revenues of $37361 compared to the $37575 earned in the unrestricted case. Thus, the cost of the initiative in the market is expected to be $214 per flight in the OD market. The total revenue change is not drastic, however, due in part to the optimal fare level of fare product 3 being close to $125 in the unrestricted optimization.

As expected, the optimal revenue result is lower than in the unrestricted case. The optimal price levels for fare products 1 and 2 are higher ($417 and $233, respectively, compared to $406 and $212) in the restricted optimization. The less-restricted fare product price level increases are expected since total demand levels are held artificially low at 125 passengers compared to the 135 passengers carried in the unrestricted case which drives fare levels up. Therefore, the increase in the fare product 3 price level has led to increases in the optimal fare levels of the other fare products. Dictating the price level of a single fare product

266

in the market not only affects that fare product but also can change the optimal price levels of the other fare products in the market. Airlines must thus be careful of the interactive effects that fare initiatives may have on the desired levels of other fare product prices.

5.6 Conclusions of the Modeling Analysis and Chapter Summary

The chapter has demonstrated the analyses available using the generalized cost model for airline fare product differentiation. Some of the relationships underlying demand, revenue, welfare, and costs were uncovered using the model as a tool for analysis. Practical applications were highlighted for the potential use of the model as a planning tool. Conclusions about the uses and applications of the generalized cost model appear in the following chapter along with directions for further research.

Chapter Six

Conclusions and Future Research Directions

The generalized cost model of airline fare product differentiation developed in this dissertation provides the first comprehensive framework that addresses how airline fare products, service bundles composed of airline prices and purchase restrictions, are viewed by consumers. The research presented here represents a departure from both the independence assumptions of existing yield management models and the techniques of stated and revealed preference measurement found in random utility models. Most notably, the incorporation of a cost for the inconvenience associated with meeting increased purchase restrictions and the interdependence of fare product price levels more accurately represents the purchase decisions confronting potential air travel consumers. The amount of revenue earned by the carrier, as well as the amounts of societal welfare and consumer surplus, vary depending upon the nature of the airline fare product "degradation" costs and of course, price levels. The types of economic efficiency achievable also vary when passenger inconvenience costs are assumed.

The main result of this research is a microeconomic model of airline fare product differentiation (and price discrimination) that incorporates the elements of product degradation, imperfect market segmentation, and capacity controls. A synthesis of the important factors associated with the prevailing fare product structure is provided by the model. In particular, the nature of the interrelationships between price levels and other airline fare product and service attributes are studied. The model can also be used to expose the relationships that exist between fare product attributes such as price level and purchase restrictions, and market performance measures like demand, revenue, and welfare.

6.1 Research Conclusions and Contributions

The conclusions presented here are subdivided into those pertaining to the generalized cost model and those involving the characterization of the airline

fare product structure that prevails in the current U.S. deregulated passenger air transportation markets. The findings are by no means mutually exclusive but the general division has been made for the sake of expository clarity. The conclusions of the airline fare product differentiation review appears first answering the question of how the prevailing fare product structure relates to issues of economic efficiency in the market.

6.1.1 Conclusions of Airline Fare Product Differentiation Review

The research presented in this dissertation examined the nature of airline fare product differentiation. The current structure of airline fare products was described in the context of the economics price discrimination literature. Contrary to the contentions of some (*e.g.* Frank, 1983), price discrimination does in fact exist in the differentiated fare product structures that dominate the airline industry today. All airline fare products offered contain second degree price discrimination or "self-selection" techniques, according to the definition used in the economics literature. Passengers self-select airline fare products according to their desired degree of travel flexibility as demonstrated by their choice of different bundles of fare product attributes.

A priori screening techniques or third degree price discrimination are also employed by airlines through such marketing promotions as student and senior discounts. Purchase restrictions and capacity controls are added to all fare products (promotional or not) available for purchase, however, and make the screened populations self-select on top of meeting any initial *a priori* screening criteria. Thus, any accurate characterization of airline fare product differentiation must model second degree price discrimination explicitly since it has been found to dominate the passenger market segmentation techniques used by airlines.

The special structure of the airline fare products existing in the U.S. domestic market today was identified and discussed in the context of product differentiation and price discrimination. An inferiority hierarchy was found to exist for the majority of fare products being sold. The hierarchical structure that currently exists on the market, the assumption of fare product independence, and the required pricing of fare products under increasing purchase restrictions

269

supports the inferiority hierarchy characterization made in the dissertation. The generalized cost modeling framework has been developed based upon this representation of the actual airline fare product structure currently facing consumers. The yield management seat allocation techniques previously developed do not explicitly consider the interdependence of the fare products based upon their price levels and purchase restrictions. Thus, the generalized cost model formulation provides insight into the interrelationships of the fare products which clearly exist but have not previously been addressed.

The dissertation also addressed questions of efficiency concerning the current practice of airline fare product differentiation. First, the question of exchange efficiency was addressed. The existence of imperfect competition, the impracticality of supply-demand equilibration, and the inability to define an unambiguous measure of short run marginal costs all prevent efficiency in exchange from occurring in airline markets. With exchange efficiency unattainable, the feasibility of achieving Pareto optimality was then addressed. Although previously hypothesized to be attainable in the current airline fare product environment, the research performed here demonstrates that costs incurred by airline passengers as a result of applied purchase restrictions make a Pareto optimal allocation of airline seats practically impossible when degradation costs are explicitly incorporated. In the context of the generalized cost model, Pareto optimality was shown to be infeasible because of the existence of degradation costs for the restricted fare products. The possibility of a more efficient allocation from an improved market segmentation method exists since a deadweight welfare loss results when any passenger incurs a cost resulting from an applied purchase restriction. Thus, two common metrics of economic efficiency have been shown to be infeasible in practice.

In spite of the Pareto and exchange inefficiencies, a different type of efficiency was shown to exist in the airline industry today : allocative efficiency. Allocative efficiency is defined as providing a limited availability service (or product) to those consumers who value it most. Any such allocation benefits society provided that resources have already been expended to produce that service. The identification of the allocative efficiency properties of airline fare product differentiation provides the first evidence of tangible benefits resulting underline{directly} from price discrimination as currently practiced by airlines.

Previous justifications for airline price discrimination have pointed to factors such as the contribution to fixed costs made by the additional passengers who would not travel unless offered a discount and the potential benefits provided by an airline earning such contributions. It has been argued that the contributions to fixed costs could be used to increase the frequency of flights in a market or decrease overall fare levels in the long term. Such benefits depend upon the airline reallocating resources to provide service improvements for consumers, which is by no means guaranteed. The allocative efficiency resulting from price discrimination, on the other hand, provides tangible benefits in terms of travel flexibility provided to passengers who place a premium on service for excess demand flights. In this vein, the differential mark-ups over marginal costs paid by these passengers can be thought of as payment for the assurance of seat availability on peak flights.

The probability of an efficient allocation was shown to be increased by a correct application of the airline yield management seat allocation techniques now in place at most carriers. The ability to secure a preferred flight itinerary with limited notice enables an airline to deliver service to those passengers who value it most on flights where demand exceeds supply. Conversely, relaxation of seat inventory controls on lower demand flights allows seats that would may have gone empty to be filled with passengers seeking discounts when flexibility-sensitive consumers face little danger of displacement. In summary, the yield management systems currently used at most major carriers have been shown to enable a desirable resource allocation in terms of allocative efficiency and to provide tangible benefits to consumers.

The requirements of a fare product structure that could support allocative efficiency in airline markets were also identified. Offering fare products with a wide range of price levels aimed at specific segments of the travel population is necessary for an efficient allocation process. This enables a finer-grained passenger segmentation. It was shown that the greater the level of passenger segmentation, the closer to allocative efficiency that an airline can come. In fact, the case of first degree price discrimination in which each consumer is identified and made to pay his exact willingness to pay guarantees allocative efficiency for the airline in the deterministic case. Although first degree price discrimination is unattainable, the closer the approximation, the more likely allocative efficiency

271

can be attained. Of course, an effective application of seat inventory control strategies must also be performed to enable an efficient resource allocation.

Efficiency discussions about the general nature of airline fare product differentiation provide insight into the characteristics of current practice in the industry. This dissertation shows that current methods, when properly applied, can result in an efficient allocation of resources particularly when demand exceeds supply. Although the equity of airline price discrimination is not clear, the more complete understanding of the efficiency characteristics of airline marketing techniques clarifies an evaluation of the equity of current industry practice. Having analyzed the underpinnings of the fare product structure, the conclusions and contributions of the generalized cost modeling framework based on the identified fare product structure appear next.

6.1.2 Conclusions of the Generalized Cost Model

Passengers neither view airline fare products as commodities nor as independent service offerings. The generalized cost model provides incorporates both the negative effects of purchase restrictions and the interrelationships of airline fare product price levels. Purchase restrictions applied to fare products decrease their value *ceteris paribus* and thus, lower the price commanded by the fare product. Such costs have not been considered previously in the yield management seat allocation models presented in the literature. The fare product interrelationships have, in general, been assumed away by viewing the different airline fare product (class) demands as independent.

As a result of the inherent dependence of airline fare products competing for space on the same aircraft, the interaction between price levels and purchase restrictions in the current practice of airline fare product differentiation should be explicitly considered in airline marketing planning models where possible. The generalized cost model of airline fare product differentiation incorporates the costs associated with the inconvenience of accepting additional restrictions associated with a given travel itinerary. The effect on demand of the increased levels of fare product restrictions as modeled here demonstrates the relationships between fare product attributes and demand from the perspective of a rational consumer.

272

Traditional airline yield management methods have failed to identify the utility that passengers purchasing lower-priced fare products have for higher-priced fare products when the price differential between the two fare products is lowered (i.e., cross elasticities have been omitted). In other words, all passengers in the market purchasing lower-priced, more-restricted fare products would purchase less-restricted fare products if the price were right. Most models predict demand levels with passenger buy down behavior implicitly assumed to exist in the population without predicting its level. An underprediction of the market revenue available to the airline has likely resulted from basing measurements of revenue potential upon observations of realized demand. Moreover, attempts at price discrimination have not been addressed in the airline yield management models. The framework of the generalized cost model incorporates the assumption that the airline can identify and segment passengers by their willingness to pay for air travel in an attempt to reconcile airline price discrimination techniques with an estimable model of airline passenger demand under fare product differentiation.

The incorporation of passenger diversion relaxes the perfect passenger identification and segmentation requirement initially assumed in the generalized cost model, allowing a characterization of airline OD market demand under imperfect price discrimination. The model can then be used to quantify the revenue dilution impacts of passenger diversion. This dissertation presents the first comprehensive discussion of the imperfections of airline price discrimination techniques with a proper behavioral motivation. The generalized cost model incorporates the passenger segmentation methods practiced by airlines with explicit consideration given to the imperfections in the segmentation techniques and their impacts.

The view of booking limits taken in the airline yield management literature has been as a method to restrict the sale of seats to independent fare products (classes) within a shared capacity to maximize expected revenue. The task of booking limits has been the protection of available seats for passengers willing to pay a higher price in exchange for the ability to obtain a seat reservation, regardless of when they book. The use of booking limits as devices used to limit the number of diverting passengers has received little attention in the literature. The generalized cost model incorporates booking limits as tools to reduce the

273

negative revenue impacts of passenger buy down behavior resulting from market segmentation imperfections.

Previously, the effects of passenger sell up have been addressed with no rigorous behavioral motivation. Passenger sell up behavior was first motivated and then incorporated into the generalized cost modeling framework developed in this dissertation. It was shown that, in order to properly address passenger sell up from a behavioral standpoint, passenger diversion must be combined with capacity constraints. Applying booking limits is shown to either induce passengers to buy a higher-priced fare product than they would prefer under unlimited fare product availability or to discourage them from traveling, depending upon their value of willingness to pay. The exact effects of applying booking limits to prevent passenger diversion are dependent upon market conditions such as the passenger booking process which dictate the amount of sell up behavior of passengers compared to rejected passenger demand.

The only passengers who are candidates to sell up to higher-priced fare products are those who have initially diverted from higher-priced fare products. This results from departures from the assumption that airlines have perfectly identified and segmented the population by their willingness to pay. After having initially been segmented, certain passengers divert to the lower-priced fare products as a result of their ability to meet the imposed purchase restrictions of the lower-priced fare products despite having a higher value of willingness to pay. Passenger diversion can only be reversed by placing capacity controls on the fare products and inducing passengers to sell up. Accordingly, the sell up potential for any OD market is exactly equal to the amount of revenue dilution that has occurred as a result of passenger diversion.

The impacts of introducing additional fare products onto the market were also explored using the generalized cost model. Under the assumption of a perfect passenger segmentation and identification, introducing an additional fare product guarantees an increase in revenues for the airline. Even in the case of perfect segmentation, however, there exist decreasing returns to the introduction of additional fare products. When the perfect segmentation assumption is relaxed, the effect of introducing additional fare products becomes indeterminate. Since passenger diversion is believed to exist in all markets, the

airline should carefully evaluate the fare products slated for introduction. The generalized cost model provides a means to evaluate proposed fare product introductions to estimate their revenue impacts under various passenger diversion and OD market price-demand relationship scenarios. The expected impacts of any changes to the prevailing fare product structure can be directly calculated from the model. The applications presented in Chapter Five provide examples of how to evaluate fare product introductions.

Airlines face several tradeoffs when evaluating their fare product structures. Among these is the tradeoff between the severity of restrictions facing consumers and the amount of passenger diversion existing in a market. The exploration of the demand stimulation versus revenue dilution tradeoff presented in this dissertation provided some insight into the choices facing airline planners designing new fare products (or modifying existing ones). Within the context of the generalized cost model, the degradation costs and passenger diversion rates embody this tradeoff. The quantification of revenue impacts of alternative fare products can be used to provide decision support to airline pricing analysts.

Another tradeoff facing airline planners is the determination of capacity for a particular market. Insight into the effects of capacity on optimal price levels and revenues can provide decision support for aircraft capacity allocation decisions. For instance, the increase in revenue potential resulting from increasing capacity may not justify the increased costs incurred by providing that additional capacity to serve a market. With greater capacity, average fare levels and thus incremental revenues, are likely to fall. If most of the passengers in a market have been accommodated with the existing capacity, allocation of additional capacity at increased cost makes little sense *ceteris paribus*. In this way, the preferred provision of capacity (and aircraft configuration) is also dependent upon the demand characteristics of the market. This dissertation has provided insight into capacity and configuration questions using the generalized cost model. Cost estimates concerning capacity decisions can be combined with the generalized cost model to improve fleet planning and aircraft scheduling decisions.

The generalized cost model has been shown to be a useful framework for analysis for airline planners, as illustrated by the examples presented in this

dissertation. Important resource allocation decisions as well as general guidelines for the pricing and design of airline fare products are provided through applications of the model. Potential uses of the generalized cost model as a planning tool have been demonstrated through the applications shown in Chapter Five and discussed here. In conclusion, the generalized cost model provides the foundation for the development of valuable decision support tools for airline planners.

6.2 Future Research Directions

The generalized cost model provides a method to evaluate the relationships that underlie passenger demand and airline revenue, in the context of the current structure of airline fare product differentiation. The model provides a method to evaluate airline pricing policies and other marketing decisions facing airline planners. Certain improvements to the basic model may provide even greater insights and are suggested below. In addition, issues surrounding the estimation of the generalized cost model are addressed. Finally, further applications and extensions of the model may unveil even more relationships important to airline pricing analysts but are left to future research efforts.

6.2.1 Extension of the Generalized Cost Model to the Dynamic Case

A static modeling methodology has been used in this research. Among the advantages of the static modeling methodology are that it is easier to estimate as a result of the lesser data requirements and, hence, greater data stability offered. The difficulties of model calibration for the multiple OD markets that face major airlines place a premium on lesser data requirements. The model also allows for comparative statics analysis to uncover the underlying relationships between price level, purchase restrictions, passenger demand and revenue. In reality, however, the booking process for airline fare products is a dynamic one which would be more appropriately addressed using a dynamic model. Such an extension would limit the types of analysis that could be performed by the model (*e.g.* no comparative statics) as well as significantly increasing the difficulty of model estimation. As a result, a dynamic formulation of the generalized cost modeling framework is left to future research.

6.2.2 Introducing Stochasticity to the Generalized Cost Model

The assumption that demand levels are deterministic for a given set of fare product price levels could be relaxed to provide a more robust model of consumer fare product purchasing behavior. While the ability of the generalized cost model of airline fare product differentiation to provide basic guidelines for the setting of price levels and the results of changing market conditions is useful, a more realistic characterization of demand may provide an improved estimate of the impacts under uncertain market conditions.

Introducing stochasticity to the price level-demand relationship occurring in the OD market under different fare product structures would require the use of stochastic functional forms of demand. The difficulties associated with stochastic mapping functions, however, are great. Techniques are available, however, which allow the incorporation of uncertainty into demand functions. For instance, the parameters associated with a chosen functional form of demand could be taken from an assumed distribution using an expected value and input to the generalized cost modeling structure. Alternatively, values could be generated randomly from an estimated distribution and used to simulate asymptotic results.

Stochasticity is more easily incorporated into the degradation costs as compared to the demand characterizations, at least in the case of the constant cost model since only a constant must be calculated with uncertainty. In the constant cost model, the single degradation cost value may be calculated from an assumed distribution of a single parameter that has resulted from an estimation based on, for example, survey results. The constant cost model could be evaluated using an expected value or simulations that randomly generate values of the degradation costs facing consumers to highlight the market variations under stochastic degradation costs.

Unfortunately, the calculation of stochastic demand or degradation cost functions provide too little benefit in the absence of a model estimated using actual data to be addressed in this dissertation. The incorporation of such stochastic elements, as well as an actual model calibration, can be achieved

through further research explorations in the area of airline fare product differentiation.

6.2.3 Estimation of Demand Functions

The generalized cost model may be calibrated based upon the assumptions made concerning the price-demand relationship in the market and the nature of degradation costs facing consumers. Once a precise form of the generalized cost model has been selected (*e.g.* the constant cost model with linear demand functions) a specific parameter estimation can be performed using presently available or, at least, collectable data sources.

The United States Department of Transportation (exhaustive 10%) coupon sample is available to all carriers (in the case of U.S. domestic airlines). The D.O.T. sample is supplemented by a coupon sample conducted by the individual carrier which reflects purchases of air travel on that carrier only. These coupon counts provide revealed preference data concerning the actual air travel choices that passengers have made. Unfortunately, the effects of passenger diversion between the fare products and the effects of booking limits have been implicitly incorporated within the fare product purchases. In addition, no information about the costs incurred by passengers associated with accepting additional bundles of purchase restrictions can be determined from the realized coupon data. Therefore, the willingness of passengers to pay for the available fare products cannot be determined from the ticket count data since the only information that is available is the realized amount that passengers have paid. Moreover, the data tell nothing about the effects passenger diversion, passenger sell up, and rejected booking requests although they are included implicitly in the data figures. Additional stated preference data must be gathered in order to estimate a demand function and degradation costs for the market.

Conducting a telephone survey to collect stated preference data about fare product choices could provide the information necessary to estimate the demand functions as well as the degradation costs for the different airline markets for which the generalized cost model is used to analyze. Although conducting a telephone survey would present many practical problems to an airline, it is likely that the wide array of airline marketing applications provided by the generalized

278

cost model would prove to be adequate reason to overcome the practical difficulties. The costs of conducting a study for every market served by even a small carrier would likely be prohibitive so the results of a survey on a limited set of markets could be used to infer the behavior of passengers in a number of markets deemed to be similar.

Unfortunately, any passenger preference survey aimed at isolating passenger sensitivities to price levels and purchase restrictions would be difficult to conduct for many reasons. For instance, deriving a measure or estimate of passengers willingness to pay is best achieved by asking questions (either directly or indirectly) about how much they are willing to pay for air travel. Respondents would have incentive to misrepresent the truth in this case for fear that a correct answer may result in higher air fares for them. Of course, such biases are inevitable when such sensitive questions are being asked. Correction of the biases would be necessary to place confidence in any demand functions calculated based upon stated preference data.

In addition, a questionnaire would likely include questions thought to be sensitive by many respondents, such as income level. It is likely that significant error will be introduced when directly asking questions concerning willingness to pay, for instance. Such response biases need to be addressed in any empirical estimation of demand incorporated into the generalized cost model. In spite of the sensitivity and error questions, however, such a questionnaire appears to be the only recourse for modeling demand under airline fare product differentiation.

Moreover, travel agents and other point of sale people have been reluctant to allow such a survey of their customers to be performed. Potential customers are likely to resent being questioned about their purchases when they are calling to book air travel. Thus, proper incentives should be offered to both perspective respondents aimed at reducing response bias as well as the point of sale contacts (*i.e.*, travel agents). Clearly, an airline conducting a survey through its own distribution channels need not worry about the latter type of compensation although additional bias correction may be warranted.

Such a survey of consumer preferences would allow the estimation of willingness to pay for all passengers within the market under study. The sensitivities of passengers to increased purchase restrictions could also be estimated from the passenger survey. Thus, an estimate of an assumed functional form of passenger degradation costs could be made. All of the information necessary to estimate the generalized cost model could be achieved from the stated preferences in a passenger survey.

6.2.4 Expanding the Scope of the Generalized Cost Model

The analysis presented in this dissertation focused on a single flight, single OD market system for a single carrier. More often in the current industry environment, several airlines offer many different flights in each OD market and each carrier operates a multiple OD market route network. More realistic applications of the model would require that more flights and OD markets be incorporated into the analysis. The extension of the generalized cost model to the multiple flight, multiple OD market would provide the framework for more powerful airline marketing planning and analysis tools. In this way, the route structure of a carrier could be explicitly considered in the analysis.

The generalized cost model of airline fare product differentiation only treated the case of monopoly. The majority of airline markets have some degree of competition. The extension of the generalized cost model to the cases of perfect and imperfect competition would provide more realistic insights into the actual pricing and fare product design questions facing airlines considering competitive effects. A model extension incorporating the interactions between carriers would provide information about the best decisions for a carrier within the existing competitive environment.

6.3 Final Conclusions

Airline revenue management consists of two interrelated functions, seat inventory control and fare product design. Although substantial theoretical work has been done on the airline seat inventory control problem, questions of fare product design and pricing have not previously been addressed. Prior

research efforts have focused on how to maximize revenues given a fixed airline fare product structure and separate, uncorrelated demands for each of the differentiated service offerings. The equally important question of how to design or modify the underlying fare product structure to improve revenues has, until this point, been ignored. Without a theoretical foundation of fare product design and pricing, only part of the airline revenue management problem has been explored.

The behavioral foundations of passenger demand in the presence of fare product differentiation and price discrimination motivated in this dissertation have provided the first comprehensive framework for analyzing airline fare product structures. A modeling framework for the heretofore unaddressed problem of pricing airline fare products under the current structure of price discrimination and product differentiation has thus been established. Moreover, many of the fundamental relationships that underlie prices and purchase restrictions in the face of a differentiated fare product structure have been identified.

The research performed in this dissertation has contributed a framework for the evaluation of fare product design based on a behavioral motivation of passenger demand under fare product differentiation and price discrimination, thereby expanding the knowledge base by addressing the full scope of the revenue management problem. This theoretical discussion of fare product design and analysis of its impacts has provided an expanded vision of the airline revenue management function explicitly considering the current techniques of fare product differentiation. The management of airline seat inventories through pricing, and its relationship to seat inventory control and aircraft scheduling, can be performed more effectively as a result of the relationships uncovered in this dissertation.

Bibliography

Aviation Daily (1991), McGraw-Hill Inc., January 29.

Bailey, E. E., D. R. Graham, and D. P. Kaplan (1985). Deregulating the Airlines: An Economic Analysis. Cambridge, MA: MIT Press.

Bailey, E. E., and J. R. Williams (1988). "Sources of Economic Rent in the Deregulated Airline Industry." Journal of Law and Economics, Vol. XXXI, April, 173-202.

Belobaba, Peter P. (1987). "Airline Seat Management - An Overview of Seat Inventory Control." Transportation Science, Vol. 21: 63-73.

Belobaba, Peter P. (1987). Air Travel Demand and Airline Seat Inventory Management. Unpublished Doctoral Dissertation. Massachusetts Institute of Technology.

Borenstein, Severin and Nancy L. Rose (1992). "Competition and Price Dispersion in the U.S. Airline Industry," NBER Working Paper.

Borenstein, Severin (1992). "The Evolution of U.S. Airline Competition." Journal of Economic Perspectives, Spring, Vol. 6, No. 2, pp. 45 - 73.

Borenstein, Severin (1991). "The Dominant-Firm Advantage in Multiproduct Industries : Evidence from the U.S. Airlines," Quarterly Journal of Economics, November, pp. 1237 - 1266.

Borenstein, Severin (1990). "Airline Mergers, Airport Dominance and Market Power," AEA Papers and Proceedings, May, pp. 400 - 404.

Borenstein, Severin (1989). "Hubs and High Fares: Airport Dominance and Market Power in the U.S. Airline Industry," University of Michigan Discussion Paper, Institute of Public Policy Studies, March.

Borenstein, Severin (1988). "Hubs and High Fares: Dominance and Market Power in the U.S. Airline Industry," Rand Journal of Economics, Autumn, 20, pp. 344-365.

Borenstein, Severin (1988). "The Competitive Advantage of a Dominant Airline," University of Michigan Discussion Paper, Institute of Public Policy Studies, August.

Borenstein, Severin (1985). "Price Discrimination in Free-Entry Markets" Rand Journal of Economics, Autumn, vol. 16, No. 3, pp. 380-397.

Brumelle, S. I. and J.I. McGill and T.H. Oum and K. Sawaki and M. W. Tretheway (1990). "Allocation of Airline Seats between Stochastically Dependent Demands," Transportation Science, Vol. 24, No. 3, August, pp. 183 - 192.

Caves, Douglas W. and Laurits R. Christensen and Michael W. Tretheway (1984). "Economies of Density Verses Economies of Scale : Why Trunk and Local Service Airline Costs Differ," Rand Journal of Economics, Vol. 15, No. 4, Winter, pp. 471 - 489.

Cross, Robert G. (1988). "The Passenger's Case for Yield Management," Presented to The Second International Airline Yield Management Conference. Washington, DC. October.

Cross, Robert G. (1986). "Strategic Selling: Yield Management Techniques to Enhance Revenue," Presented to Shearson-Lehman Brothers 1986 Airline Industry Seminar. Key Largo, FL. February.

Elkins, Steven (1986). "Summary of Basic North American Pricing/Inventory Management Policies," Northwest Airlines interoffice memo, October 13.

Douglas, George W., and James C. Miller (1974). Economic Regulation of Domestic Air Transport. Washington: Brookings Institution.

Frank, Robert H. (1983). "When are Price Differentials Discriminatory?" Journal of Policy Analysis and Management, Vol. 2, No. 2, pp. 238 - 255.

Fuchs, S. (1987). "Managing the Seat Auction," Airline Business. July, 40-44.

Glover, Fred, Randy Glover, Joe Lorenzo, and Claude McMillan (1982). "The Passenger-Mix Problem in Scheduled Airlines," Interfaces. 12(3): 73-79.

Graham, David R., Daniel P. Kaplan, and David Sibley (1983). "Efficiency and Competition in the Airline Industry," Bell Journal of Economics, 14:118-136.

Hirsch, James S. (1993). "American Air's Flyers Won't Get Fare Warnings," Wall Street Journal, p. B1, February 25.

Henderson, Danna K. (1986) "Carriers Report Significant Increases in Passenger Handling Costs," Air Transport World, September.

Holmes, Thomas J. (1989). "The Effects of Third-Degree Price Discrimination in Oligopoly," American Economic Review, March, Vol. 79, No. 1, pp. 244 - 250.

Jara-Diaz, Sergio and Clifford Winston (1981). "Multiproduct Transportation Cost Functions: Scale and Scope in Railway Operations." in Eighth European Association for Research in Industrial Economics, Vol. 1, eds. Nicklaus Blattner et al. U. of Basel, 1981, pp. 437-69.

Kahn, Alfred E. (1988). The Economics of Regulation : Principles and Institutions, The MIT Press, Cambridge, Massachusetts.

Keeler, Theodore E. (1989). "Airline Deregulation and Market Performance: The Economic Basis for Regulatory Reform and Lessons from The U.S. Experience," Presented at the ESRC Seminar on Privatization and Deregulation of European Airlines, Oxford, U.K., September 8.

Keeler, Theodore E. (1978). "Domestic Airline Regulation: An Economic Evaluation," in L. Weiss and M. Klass, eds., Study on Federal Regulation, U.S. Senate, Committee on Governmental Affairs, Vol. VII, Washington: U.S. Government Printing Office, 77-160.

Keeler, Theodore E. (1972). "Airline Regulation and Market Performance," Bell Journal of Economics, 3, Autumn, 399-424.

Keeler, Theodore E., and Michael B. Abrahams (1981). "Market Structure, Pricing, and Service Quality in the Airline Industry under Deregulation," in W. Sichel and T. Gies, Applications of Economic Principles in Public Utility Industries. Ann Arbor: Michigan Business Studies, Vol. 2, No. 3, 103-145.

Koran, Donald W. (1983). "The Welfare Effects of Airline Fare Deregulation in the United States, Journal of Transport Economics and Policy. May, pp. 177- 189.

Lee, Anthony O. (1990). Airlines Reservations Forecasting: Probabilistic and Statistical Models of the Booking Process. MIT Flight Transportation Laboratory Report 1990-5. Massachusetts Institute of Technology.

Levine, Michael E. (1987). "Airline Competition in Deregulated Markets: Theory, Firm Strategy, and Public Policy," Yale Journal on Regulation, 4, Spring, 393-494.

Morrison, Steven, and Clifford Winston (1990). "The Dynamics of Airline Pricing and Competition," AEA Papers and Proceedings, May, pp. 389 - 393.

284

Morrison, Steven, and Clifford Winston (1986). The Economic Effects of Airline Deregulation. Washington: Brookings Institution.

Morrison, Steven, and Clifford Winston (1987). "Empirical Implications and Tests of the Contestability Hypothesis," Journal of Law and Economics, Vol. XXX, April.

Morrison, Steven, and Clifford Winston (1989). "Enhancing the Performance of the Deregulated Air Transportation System," Brookings Papers on Economic Activity, Microeconomic Issue, 1989, 61-112.

O'Brian, Bridget (1992). "AMR's Bid for Simpler Fares Takes Off," Wall Street Journal, pp. B1, B6, April 10.

Official Airline Guide: Worldwide Edition (1993). Official Airline Guides, Inc., Oakbrook, Illinois, Vol. 16, No. 8, August.

Pfeifer, Phillip E. (1989). "The Airline Discount Fare Allocation Problem," Decision Sciences, Vol. 20, pp. 149 - 157.

Pigou, A.C. (1920). The Economics of Welfare, fourth edition, Macmillan, London.

Salop, Steven C. (1979). "Monopolistic Competition with Outside Goods," Bell Journal of Economics. 10, 141-156.

Salop, Steven C., and David T. Scheffman (1983). "Raising Rivals' Costs," AER Papers and Proceedings, May, Vol. 73, No. 2, 267-271.

Schmalensee, Richard (1981). "Output and Welfare Effects of Third Degree Price Discrimination," American Economic Review, March, Vol. 71, pp. 242-47.

Schweiterman, Joseph P. (1985). "Fare is Fair in Airline Deregulation," AEI Journal on Government and Society, pp. 32-38.

Simpson Robert W. (1982) Notes for Air Transportation Economics, Massachusetts Institute of Technology, Cambridge, Massachusetts.

Smith, Barry C. and Crystal Penn (1988). "Analysis of Alternate Origin-Destination Control Strategies," Agifors Symposium, October.

Tirole, Jean (1988). The Theory of Industrial Organization, The MIT Press, Cambridge, Massachusetts

United States General Accounting Office (1990). "Airline Competition: Higher Fares and Reduced Competition at Concentrated Airports," July.

Wang, Judy S. and Ann Friedlaender (1981). "Mergers, Competition and Monopoly in the Regulated Trucking Industry." MIT Working Paper No. 289, 1981.

Weatherford, Lawrence R. and Samuel E. Bodily (1992). "A Taxonomy and Research Overview of Perishable -Asset Revenue Management : Yield Management, Overbooking and Pricing," Operations Research, Vol. 40, No. 5, September-October, pp. 831 - 844.

Weatherford, Lawrence R. (1991). "Perishable -Asset Revenue Management in General Business Situations." Unpublished Ph.D. Dissertation, Darden Graduate School of Business Administration, University of Virginia, Charlottesville.

Wells, Alexander (1984). Air Transportation: A Management Perspective, Wadsworth Publishing.

Williamson, Elizabeth L. (1992). Airline Origin-Destination Seat Inventory Control. MIT Flight Transportation Laboratory Report 1992-3. Massachusetts Institute of Technology.

Winston, Clifford (1985). "Conceptual Developments in the Economics of Transportation: An Interpretive Survey." Journal of Economic Literature, Vol. 23, pp. 57-94, March.