FLIGHT TRANSPORTATION LABORATORY
REPORT R95-8

THE VALUE OF REVENUE MANAGEMENT INNOVATION IN A COMPETITIVE AIRLINE INDUSTRY

JOHN L. WILSON

MAY 1995

The Value of Revenue Management Innovation in a Competitive Airline Industry

by

John L. Wilson

B.A., Economics
Cornell University, 1993

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

MASTER OF SCIENCE IN TRANSPORTATION

at the

Massachusetts Institute of Technology

May 1995

Signature of Author . . . .
Department of Civil and Environmental Engineering
26th May, 1995

Certified by . . . .
Peter P. Belobaba
Assistant Professor of Aeronautics and Astronautics
Thesis Advisor

Accepted by . . . .
Joseph M. Sussman
Chairman, Departmental Committee on Graduate Studies

The Value of Revenue Management Innovation in a Competitive Airline Industry

by

John L. Wilson

Submitted to the Department of Civil and Environmental Engineering on 26 May, 1995, in Partial Fulfillment of the Requirements for the Degree of Master of Science in Transportation

Abstract

The value of revenue management to the airlines has been amply demonstrated, both by industry experience and in simulation studies of the reservation process. However, there have been no attempts to determine if the benefits of seat inventory control are specific to the competitive market setting in which it is instituted. Furthermore, previous theoretical research has not considered whether interactions of the control methods of rival carriers may affect both the total revenue improvement attainable in the market and how these gains are shared by individual carriers.

This thesis uses a modeling approach created by the Boeing Commercial Airplane Group to address these questions. The Passenger Origin / Destination Simulator combines a demand model framework with a set of routines which implement the basic forecasting and inventory control functions of airline revenue management. The Boeing simulation system is implemented in the thesis through the development of a generic fare product structure, passenger behavioral attributes, and assumptions on demand composition.

Within this Operational Competitive Simulation Environment, the effect of competition on the value of revenue management is explored in simulation experiments with three classes of scenarios. In these scenarios, the gains from leg-based inventory control are assessed under varying competitive conditions, including the magnitude and distribution of passenger demand, carrier frequency share, relative departure timing, and route network design. From the simulation results, it is apparent first that inventory control innovation always improves carrier and aggregate market revenues—competition along the dimension of revenue management is not a zero-sum game from the carriers' perspective. A signal contribution of the experimental research has been the finding that the revenue dividend from inventory control derives largely from the sale of higher-yield fare products to leisure and discretionary business passengers, and only marginally at the expense of those carriers with a less advanced control capability. Additionally, the relative changes in revenue due to different control methods are indeed influenced by other competitive forces operative in the market.

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

Acknowledgments

During my first semester at MIT, Peter Belobaba instilled in me his enthusiasm for revenue management and the economics of the airlines. Several months later, the timely completion of this thesis has been made possible by his confidence and guidance in the final weeks of writing. From start to finish, I count myself fortunate to have been his student and colleague.

Craig Hopperstad might first be recognized as the creator of the simulation system which is the foundation of this work. However, his patient explanations of the model's intricacies, and comments on an early draft of my motivational arguments, have also strengthened the logical development of the thesis considerably. Craig's dedication and genial nature, which shined through even the darkest days of debugging, have been great assets in our joint model development and implementation efforts.

Daniel Skwarek deserves no small credit for bearing with my eccentricities in the production of many of the supporting graphs contained herein. He joined the project at a critical phase of the scenario development, and contributions in this area, as well as his review of my background sections, are gratefully acknowledged.

It has been a pleasure working with Peter, Craig, and Daniel over the past year. Given their collective resolve and capability, I leave assured of the future success of this collaborative research project.

Table of Contents

List of Figures

List of Tables

1 Introduction

1.1 Setting, Purpose, and Motivation

As the concept applies to the modern airline industry, revenue management (RM) represents nothing more than the profitable allocation of scarce resources—airplane seats, here—in the market for air transportation services. While this broad objective is simply stated and indeed can be identified as a tenet of elementary microeconomics applicable to any business, its pursuit in the airline context has stirred an unusually strong degree of interest, both academic and operational, among practitioners and theoreticians viewing the problem from the carriers' perspective. Counterbalancing this enthusiasm is the outcry from disillusioned passengers who understandably object to feeling manipulated and must devote considerable resources of their own in navigating an increasingly chaotic array of fares.

RM is also known as *yield* management, since it seeks to alter the mix of yields, or per passenger-mile revenues, which would otherwise be collected in an "unmanaged" setting, where reservations are accepted on a first come, first served basis. If, on a given flight with one unoccupied seat remaining, an airline has the choice between selling a full-fare ticket to a corporate executive or a deeply-discounted ticket to a vacationer, it will naturally try to attract the former high-yield clientele. However, the life of a revenue manager is complicated by two fundamental characteristics of air travel: demand from each passenger segment is random (i.e., it is not known for certain whether the hypothetical executive will in fact attempt to find space on a given departure); and, as time-sensitive business travelers expect last-minute availability while holiday-makers plan their itineraries further in advance, lower-yield traffic typically books passage first. Herein lies the central dilemma of RM efforts—should a seat be made available to a leisure traveler inquiring two weeks before departure, *potentially* closing out a higher-yield business customer who would only state his intent to travel shortly before departure?

In the parlance of game theory, the "payoff matrix" resembles Figure 1.1. Alas, there is no "dominant strategy" which may be played in these circumstances. Specifically, accepting the vacationer (at a fare of $100) is optimal under only one of the two possible states of the world—when the executive (who would be charged $300) fails to materialize. Similarly and symmetrically, holding out for the executive by dismissing the sure $100 may backfire. Fortunately, real world RM analysts have a host of armaments at their disposal—sophisticated forecasters, optimizers, and evaluators of historical performance ("What worked in the past?")—to encourage appropriate action. Then again, this stylized scenario admittedly neglects many of the complications which plague true RM decisions, including the diversity of fare products, route network effects, and forecasting difficulties. Chapter 2 surveys the full extent of the RM problem and its proposed solutions.

	Subject to Probablistic Demand Variation	
Airline Decision	Executive Later Requests Reservation for Fare of $300	Executive Chooses to Teleconference
Vacationer Accommodated for Fare of $100	*third best outcome* receive $100 of potential $300	*second best outcome* receive $100 of potential $100
Vacationer Refused	*best outcome* receive $300 of potential $300	*worst outcome* no revenue received of potential $100

Figure 1.1. Airline Payoff Matrix for the Revenue Management Decision

Yet, after reviewing just these barest essentials of the RM problem, the title of this thesis can be recalled to introduce its purpose: What is the value of adopting RM *in a competitive setting?* Consider three eras in recent airline history:

- *No RM*: It may be difficult to imagine a world before RM, but in fact, prior to the Airline Deregulation Act of 1978, airlines had little of the incentive and none of the means to practice discriminatory pricing—even if such behavior had not been explicitly declared illegal. Inventory management had barely entered the realm of speculation.
- *RM as innovation*: In the interim, the adoption of modern RM heuristics and hardware has been neither instantaneous nor uniform. Developments pioneered by those major US carriers blessed with stable research and development budgets have only gradually filtered into common industry usage.
- *RM as standard*: Today, fifteen years after its potential first became apparent, most significant domestic airlines have achieved at least a basic proficiency in RM.

This thesis endeavors to determine both the competitive advantage won by individual rival airlines wielding RM control and the aggregate benefit granted to the industry at large—in each of the latter two evolutionary phases—relative to the first. The inquiry naturally inspires many interrelated provocative questions:

- In the cases where everyone operates on a "level playing field", does the overall industry fare better when all carriers have RM or when none do?
- How does the revenue reward for a single revenue-managed airline vary when:
  - no market competition exists?
  - the competition has no RM?
  - the competition has an equivalent RM capability?
- Which aspects of passenger demand significantly affect the reward? Can other airline decisions concerning scheduling and route structure be coordinated to maximize RM effectiveness?
- Who provides the gains which accrue to successful RM: passengers or rival carriers?
- For a given set of demand conditions, is inter-carrier competition for market revenue a "zero-sum game," in the same sense that a market share gain for one airline necessarily comes at the expense of another's loss? Or, does the entire revenue "pie" grow larger under certain combinations of RM policies?

Airlines closely guard their RM technology and related performance measures, precluding direct testing of the hypotheses suggested above. Therefore, this investigation must be performed through simulation.

1.2 Outline

Chapter 2 reviews the essence of RM theory and practice. Virtually all of the work performed by both the academic and airline communities in the last few years consider the promise and pitfalls of what would be a sea change in current RM practice—the evolution from today's leg-based control systems to those offering some level of network or origin-destination (OD) optimization. Two approaches towards this end have emerged. The airlines, conscious of their enormous sunk investment in the existing reservations infrastructure, have endeavored to work within a leg-based architecture, modifying leg-based RM structures to reap an incremental revenue benefit, while others less concerned with the details of implementation and operations, propose an entirely new theoretical control framework. While these relatively exotic developments undoubtedly presage the future of successful RM, it is worth first reflecting on the significance of the competitive advantage yielded by basic, universally-accepted RM techniques. Embarking on a course likely to be characterized by diminishing returns and accelerating expense, it must be determined how the gains of RM are generally distributed among individual participating carriers, and whether the industry at large enjoys a net benefit.

Chapter 3 introduces the simulation setting in which the experiments will occur—the Passenger Origin-Destination Simulator (PODS) created by researchers at the Boeing Commercial Airplane Group. After an overview of the PODS system architecture, outlines are presented for the three major submodels—stochastic demand generation, passenger assignment, and inventory control and forecasting—and the flow from user inputs to the output of performance measures are traced.

Further definition of the PODS framework has produced the Operational Competitive Simulation Environment (OCSE). The OCSE assumptions, including the fare structure and behavioral characteristics, follow in the first half of Chapter 4. The remainder of the chapter considers the sensitivity of these settings with simulation trials under a noncompetitive schedule. The tested dimensions include demand magnitude and variability, traffic composition, and valuation of ticket restrictions.

Chapter 5 presents the descriptions and results for simulations of three classes of competitive market scenarios. The scenarios test the effect of demand magnitude, relative timing and frequency of departures, and path quality, under various RM method combinations. Evidence for the thesis that the degree and nature of competition affects the value of RM appears in the form of graphs and tables, illustrating revenue and traffic statistics for the total market and how benefits are shared among individual carriers.

Chapter 6 first draws together the major findings and their implied consequences for the economic rewards possible with a basic inventory control capability in different competitive settings. The next steps for the PODS / OCSE modeling approach are then described. Though the scope of this thesis does not allow consideration of next generation OD-control systems, the framework devised here will enable similar analyses of alternative optimization and forecasting methods which can be translated for use within PODS. Even within the existing scope of RM alternatives, significant enhancements can be made to the current reservation process model. Finally, more detailed competitive scenarios may be explored to gain further insight into the relationship between the value of RM and the market environment in which control is implemented.

2 Revenue Management Principles and Research Directions

The yield-management problem is best described as a nonlinear, stochastic, mixed-integer mathematical program that requires data, such as passenger demand, cancellations, and other estimates of passenger behavior, that are subject to frequent changes. To solve the system-wide yield-management problem would require approximately 250 million decision variables. (Smith, *et al.*, 1992)[1]

2.1 Revenue Management and the Provision of Airline Services

In the airline production process, RM decision variables are the most tactical in nature. A stylized representation of the typical carrier's operation would begin with aircraft acquisition and the attendant financing and budgeting. A route network is next developed on the basis of economic OD market analyses. In the third stage of fleet and crew assignment, the carrier's human and capital resources are associated with the proposed routes, negotiating both physical constraints and those enforced by labor contract. The service schedule which emerges can be thought of as the airline's "supply." Passengers are now confronted with a slate of "goods" which will give rise to some level of demand. At this point, RM enters the picture, in its role of achieving a profitable equilibrium between the established schedule and the public's derived travel demand.

The corporate hierarchies of most airlines contain a subdivision placing two offices within the province of RM: pricing and seat allocation. The first responsibility, despite its apparent simplicity, needs some elaboration. Fares are set for the various itineraries available in each OD market and for the service options offered within each routing. All else equal a seat on a nonstop departure typically carries a price *lower* than a through or connecting ticket for travel between the same cities. Given that the average passenger would prefer to minimize trip time by choosing the nonstop, such a fare system violates the principle of, "charging what the market will bear." However, this seeming irrationality can be justified by the congested state of many spoke flights, the components of the modern hub carrier's connecting options. To preserve space for passengers traveling in markets only accessible through the hub, airlines adopt the counterintuitive pricing policy

[1] p. 9.

to draw traffic to their nonstop service when it is available. In a more straightforward example, a space in the luxuriously appointed first class cabin would be more expensive than a coach seat in the same plane.

While these distinctions are clear and logical, they appear to fall short of addressing the overriding complaint of modern airline travel: how can two neighboring economy class passengers, enjoying identical levels of service, be charged wildly varying fares? The common resolution to this seeming paradox denies the premise. While they may experience the same in-flight amenities, other service dimensions, connected to the ticket restrictions which must be met at the time of reservation, should not be overlooked. These include advance purchase (e.g. ticket is last available at 14 days prior to departure), rules on permitted trip characteristics (e.g. minimum / maximum time at destination; Saturday night stay), and limits on refundability.

Far from arbitrary, such carefully constructed "fences" enable imperfect price discrimination, and have been an integral feature of airline economics since deregulation. Packages of restrictions, together with the monetary fare level, constitute *fare products* coded under an arcane scheme in which the letters "Y" and "Q" can signify the two ends (high fare / unrestricted and low fare / highly restricted, respectively) of the coach class spectrum.[2] Under a more generalized notion of service quality then, the most desirable, flexible Y fare product fittingly carries the highest price tag. Nevertheless, appraisal of restriction burden is in the eye of the beholder, and Chapter 3 returns to the subject of restriction cost valuation in the introduction of the behavioral parameters included in OCSE.

The second mission of RM is to control the number of seats made available to inquiring passengers throughout the booking process with the goal of revenue maximization.[3] For any flight, this inventory management process starts with the capacity of the aircraft assigned to the route. This figure is then adjusted upward to compensate for passengers who, for some reason, fail to keep their reservation. Systematic overbooking, in recognition of the inevitable but variable "no-show" phenomenon, has a long and

---

[2]There is, however, no industry-wide naming convention. The four fare classes adopted in the scenario analyses of Chapters 4 and 5, representative of the current industry hierarchy in fares and restrictions, have been given the generic codes, Y-M-B-Q.

[3]Although other objective function terms are conceivable (including profit, yield, and load factor), gross revenues generally draw the most attention and will be the focus of comparisons made later in the thesis.

colorful history reaching well into the regulated period, and debate continues today over appropriate overbooking factors which balance the benefits of higher loads with the costs (if any) of denied boardings.

However, in contrast with practice of twenty years ago, inventory management today only begins with the establishment of the overall authorized booking level. Far more effort must be devoted to ascertaining the optimal distribution of this total limit over each of the fare classes created by the pricing department. Indeed, the dynamic fare environment that has characterized air travel since deregulation is due almost entirely to the sophistication of seat allocation algorithms, and only incidentally to the absolute fare product price structure. The reason for this lies at the heart of airline competitive strategy. The published price levels appearing in bold print in newspaper sale announcements for any one carrier permit no response other than immediate matching on the part of all major rivals, at least in all shared markets.[4] The battle is really won or lost in the fine print, so to speak, where customers learn that capacity constraints, black-out dates, and other caveats apply to the advertised fares. Airlines skilled in inventory management minimize the revenue impact of sales initiated by low-cost entrants and other majors in financial straits, while maintaining the essential outward appearance of a sale. The accuracy of the initial seat allocations and adjustments of these limits throughout the booking process decide the short-term performance of an airline, given the operating environment created by prior market entry, fleet assignment, and schedule positioning decisions. Understandably then, the industry has continuously sought to refine the forecasting and optimization components of its RM technology, while realizing that the immensity and near real-time nature of the problem preclude a search for theoretically optimal answers.

2.2 Revenue Management Theory

2.2.1 The Formal Revenue Management Problem

Pioneered in the airline context, the principles of yield management have in recent years been extended to other associated industries, such as hotels and different transportation modes. Through a comprehensive classification scheme, Weatherford (1992) characterizes the general class of RM problems, as well as those distinguishing attributes which describe specific application settings.

[4]Exceptions to this rule find majors selectively ignoring low-cost carrier discounts when supported by a significant hub presence ("home town image") or clear schedule advantage in a market.

All environments which stand to benefit from RM analysis share three basic properties: product perishability, limited capacity, and the possibility of market segmentation. The finite lifetime of the airline product—a seat on a particular departure—clearly meets the first criterion. Seats cannot simply be stored in times of low demand and released just-in-time.

Furthermore, within the short-term planning horizon relevant to revenue managers, aircraft capacity on each route must be taken as fixed.[5] Fleet assignment, occurring perhaps every several months, cannot immediately remedy flight schedules in which demand consistently exceeds aircraft size for some departures. Even allowing for aircraft exchanges, it may not make sense to devote a model large enough to accommodate historical demand, given larger network considerations. Certainly no airline has the resources to eliminate capacity constraints in its system, nor would the cost side of aircraft operations render it profitable to do so. Inevitably then, it is RM's responsibility to control the allocation of this rigid supply.

Finally, RM will be possible only if passenger segments of varying price sensitivities can be identified and their preferences implicitly captured in the line of fare products. In general microeconomic theory terms, a fare hierarchy must be established in which the cross-price[6] (between product) elasticities of adjacent products is low relative to the own-price (within product) elasticities. The most common differentiating mechanism concerns time-of-purchase. If the trade-off between dollar price and flexibility (measured by the severity of the advance purchase restriction) inherent in the fare structure is deemed acceptable by a sizable group of customers, incremental sales can be made under this fare class to improve capacity utilization and total receipts. For the low variable seat cost (consisting of minor passenger handling and in-flight nondurable expenses), often found juxtaposed with a high fixed capacity cost, makes it worthwhile to accept a fare below even the deepest excursion discount, rather than depart with the seat empty. Of course, this argument holds only on a marginal basis. Filling systemwide capacity with Q class traffic, as charter outfits do, would not cover the overhead borne by the majors. As

---

[5]Berge (1993) suggests that under "demand driven dispatch," dynamic aircraft assignment via heuristic algorithms may partially correct the demand/capacity imbalance. Nevertheless, some degree of capacity "lumpiness" (discrete seat quantities) and "stickiness" (lag, or other resistance, to change over time) is inescapable.

[6]Where price is taken as a generalized term, incorporating the degradation costs of the attached restrictions.

Belobaba (1987) argues, a heterogeneous marketplace and the pricing freedom encouraged by deregulation allow for a scale and quality of service which would not be achievable under a uniform, nondiscriminatory fare system.

Several additional elements of the taxonomy define the RM task facing the airlines. First, seat resources must be managed in discrete units. Several seat allocation approaches rely on mathematical programming techniques whose complexity grows considerably with these integrality constraints.

The proliferation of fare products, fielded in response to special competitor marketing initiatives, also contributes to the magnitude of the problem. As many as ten inventory classes on each flight leg must be jointly managed. As the existence of ten stable and distinct passenger categories seems unlikely, predicted class demand arrival rates and totals are subject to large errors.

Regarding the actual execution of RM assignments, the intimately related seat allocation and pricing responsibilities are typically performed sequentially. Booking limits for each class then, are made conditional on the prevailing fare levels. Forcing inventory managers to work within a predetermined pricing structure sacrifices degrees of freedom in the combined revenue maximization effort and necessarily leads to a second-best outcome. This said, airlines cannot generally achieve the ideal theoretical solution—simultaneously optimizing price and supply—within operational constraints. Additionally, the two revenue determinants do ultimately interact in feedback loops of varying strengths and frequencies. Whether successive iterations of the RM process have time to converge on the simultaneous result before the next sale crops up and rearranges the landscape, is debatable.

Common airline passenger behavior poses other problems by increasing uncertainty. At the highest level, practical RM techniques must reflect the stochastic nature of demand. Many further suppose that reservations arrive independently among fare classes. This is not to say that all passenger segments follow completely random arrival rates over the booking horizon. Indeed, willingness to pay, as a rule, rises as the departure date approaches. Nor are passengers obliged to remain in their "proper" market segment—i.e., the one they (or their historical counterparts) were associated with in the forecasting process. Floating price- and time-sensitivities leave open the possibility for the double-edged sword of *diversion*. If a traditional full-fare payer finds it possible to break ranks and meet discount restrictions, unwelcome *dilution* occurs. However, seat managers also cannot foresee the consequences of turning away an authentic discount request—the refused passenger may: a) be persuaded to "buy up" to a higher fare class on the same flight; b) find the initial fare available on a succeeding flight; or c) exit the airline's reservation process entirely, either by flying a rival carrier or postponing travel plans. Options a) and b) still affect inventories, and a) carries a positive revenue effect.

Group reservations—in which multi-sized blocks of customers, having correlated preferences, must be collectively accommodated or denied—are another hindrance to the development of smooth booking profiles. Rounding out the revenue manager's plight is a flight-specific cancellation / no-show factor which leaves the ultimate success of the RM mission undecided until the very day of departure.

2.2.2 Optimization Scope

RM control methods can be broken into two classes by the scope of their optimization routines: *leg-based* systems and *network* or *OD-based* systems. To explain the difference between these two categories, it is helpful to introduce a few terms commonly used to describe air transportation operations.

Confusion often arises in these descriptions because, as in many transportation system contexts, there are really two network characterizations which relate to air travel. One network illustrates how an airline *routes* its aircraft fleet among the cities it serves. Given a network of cities (the nodes of a generic network plan), the carrier determines which city pairs will be directly connected by nonstop departures (links, or arcs). These adjacent city pairs are known as *sectors*. The *schedule* for an individual aircraft describes the sequence of sectors traversed over a schedule cycle of one or more days. Each nonstop component of the schedule is referred to as a *flight leg*. The distinction between a flight leg and a sector is that the former is associated with particular departure and arrival times.

To the passenger, much of the network of aircraft routings is transparent. With a derived demand for air travel, passengers wish only to reach a certain destination from their origin, in a specified window of time. They accomplish this objective by piecing together a *path* or *itinerary* from the flight legs offered by carriers serving the city pair. Thus, the aircraft and passenger networks overlap but do not necessarily coincide. A passenger may remain on one plane for the duration of the trip, on nonstop or continuing / through itineraries. Alternatively, in the hub route networks which have proliferated in the post-deregulation environment, a path may be composed of several flight legs on multiple aircraft.

The operating efficiencies of hub networks, which have been well demonstrated in theory and practice, are the result of funneling passengers from multiple OD markets to spoke flight legs emanating from the hub. However, when passengers traveling in OD markets with different fares share the capacity of a single flight leg, RM revenue maximization becomes much more difficult, as will be seen below.

Leg-based control methods optimize inventories on each flight leg of a network independently, tacitly ignoring matters of flight connectivity and passenger flows over multi-path legs. *OD-based optimization* offers a partial-to-full accounting of hub network effects, depending on the formulation. Bridging these two classes, the stopgap control approach known as *virtual nesting* incorporates local leg interactions by selective itinerary and fare product recategorization within an airline's internal reservation control structure (see Section 2.2.2.2).

Although pure leg-based schemes have come into disfavor in the last few years, even the largest carriers continue to rely, at least partially, on such systems and the majority of the world's airlines use nothing else. Given their enduring popularity, and in preparation for the central simulation study of this thesis, leg-based models will be the focus of the discussion below. Readers wishing further details on seat inventory algorithms may consult Williamson's (1992) critical review which runs from the first formal optimization rules of the early 1970s to current OD proposals. Before turning to these alternative scopes, common aspects of both the forecasting process and the booking controls themselves are explored.

It has not yet been explicitly stated how inventory managers arrive at *booking limits* for a certain flight. Although the details vary from one carrier to the next, forecasting always begins with the establishment and population of an historical traffic database. Complete booking profiles for previous departures can be drawn from the database, represented by cumulative booking curves similar to the ones illustrated in Figure 2.1. These may be aggregated in some fashion (as by length of haul, below) or applied to individual flight numbers. Day of week is a common forecast criterion—Friday departures are best predicted with prior Friday results. Corrections must also be made for seasonal trends and other unusual demand circumstances. Should the competitive stability of the flight's OD markets be upset by new schedule or fare initiatives, forecasting becomes more of a subjective exercise dependent on elasticities derived from the observed effects of similar disturbances in the past.

Beyond these basic storage and retrieval functions, RM decision support capabilities vary considerably (Belobaba, 1992b). Some systems merely track bookings for future departures, alerting the human analyst to those demanding special attention. When thresholds developed from historical trends are violated, the affected seat inventories may be adjusted manually, perhaps with guidance on a general course of action from the monitoring computer. However, analyst response may be subjective, inconsistent, and even detrimental in some cases. Exponential improvements in processing speed, have precipitated the progression to automated booking limit systems. Automated managers allow systematic implementation and subsequent evaluation of RM strategies, consolidating the experience of many individual human analysts.

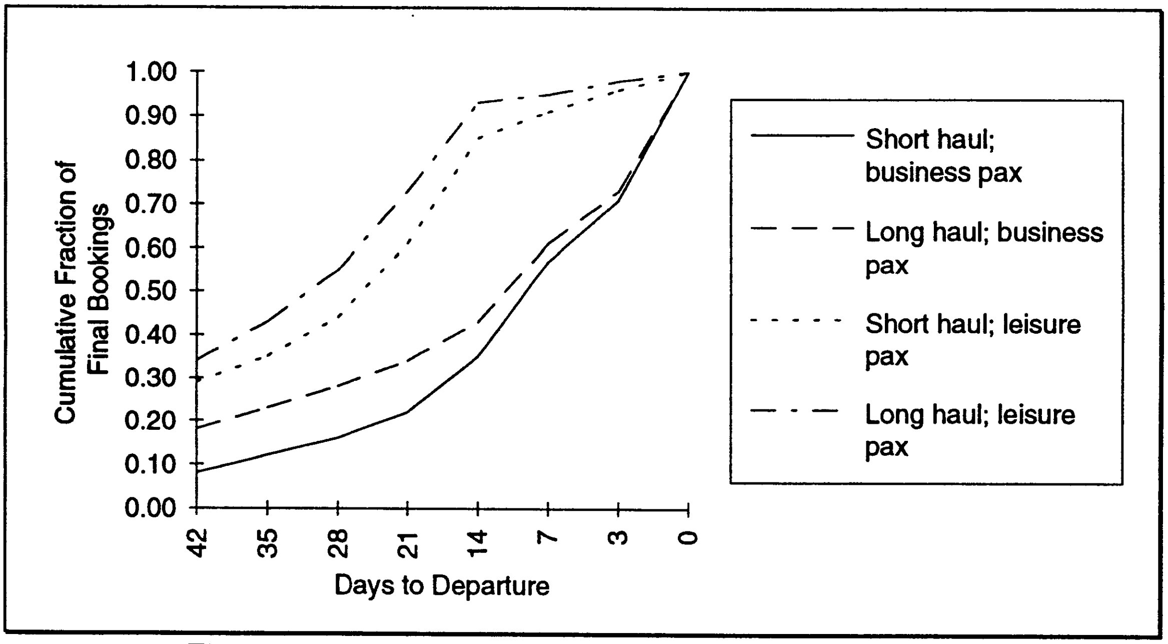

Figure 2.1. Booking Patterns by Market and Passenger Types

Automation has also expanded the opportunities to fine-tune inventory allocations over time. The entire booking horizon, beginning up to one year before departure, is divided into *booking periods* or *time frames* of various lengths (shortening as the day of departure approaches). Limits on the number of seats available to each fare class are established at discrete moments, known as "revision points," at the start of each booking period. Examination of the partially completed booking record as of an intermediate revision point, and comparison with the same period in prior departures, may lead to adoption of a modified rule for inventory control decisions. It is intuitively clear that such *dynamic* updating of the reservation controls will outperform *static* systems, which do not deviate from the allocation prescribed 330 days before departure—one necessarily based entirely on historical data. Computerized systems have made the use of frequent revision points (i.e., daily, close to departure) commonplace. Repeated iterations of an efficient optimization heuristic show an insignificant performance shortfall against the prohibitively expensive ideal.

One recurring concern in forecasting methodology is the appropriate handling of situations in which the number of accepted bookings matches the absolute capacity limit. Such cases leave uncertain the magnitude of the true, *unconstrained* demand which would have been observed if there were no ceiling on capacity. A carrier must assess unconstrained demand to decide whether capacity expansion, in the form of an extra flight or larger aircraft, would be beneficial at the prevailing fares. Unconstrained demand information is also required to determine whether the carrier has allocated enough seats to the higher fare classes on previous departures. It is possible that the capacity constraint did not in fact cause the refusal of any booking requests—that is, unconstrained demand just happened to match capacity. Use of the raw booking data would then accurately capture the true demand conditions. It is more likely, though, that capacity was indeed a binding constraint, and unconstrained demand exceeded the level recorded in the reservations system. If the constrained data were applied, the resulting forecasts would be biased downward. Omission of the problematic observations would produce an even more distorted result.

Hence airlines have sought a suitable *detruncation* mechanism. Conceptually, the true demand must be extrapolated from the shape of the early booking profile. When the slope of the cumulative booking count is steep, the flight or an individual fare class will "close" early (upon reaching the capacity or booking limit). Between the time of closure and the day of departure, many requests might have been received and turned away. In contrast, if the flight closes just before the day of departure, the unconstrained demand probably did not exceed the capacity limit by very much. The simplest and most popular detruncation approach scales up the closed bookings in a linear fashion from the point the flight closed (Section 3.4 describes PODS implementation of the detruncation function).

As to the forecasting technique itself, the most common approach, known as the *pick-up* model, rests on the belief that current bookings have no bearing on later reservation activity. In pick-up analysis, the only quantity to be determined is the mean number of bookings to expect in a future period. An alternative to pick-up, adapted from other fields which rely on forecasting, is linear regression. Applied to the RM context, a simple regressed relationship takes the form:[7]

$$BTC = a + b(BIH) \tag{2.1}$$

where

BTC is the number of bookings to come in the future, the quantity to be explained,
BIH is the number of bookings received to date ("in hand"),
a, b are the intercept and slope terms, respectively, regressed from records of past bookings.

Regression offers insight into the pattern of bookings over time through its slope term. Correlation across booking periods, whether positive or negative, would be indicated by a significant value for b of the corresponding sign. Pick-up can be seen as a special case of regression in which the slope term is forced to be zero. Given a statistically significant positive slope term regression would be expected to outperform pick-up. However, regression can also "detect" nonexistent booking trends, leading to a serious misallocation of capacity.

A final important distinction concerns what Weatherford refers to as the "asset control mechanism" in his generic taxonomy. *Nested* systems set the booking limit for the highest fare class as the total authorized booking limit for the flight—aircraft capacity plus an overbooking factor. The limit for a lower class is found by subtracting the *protection level*[8] for the next higher class from the limit for the higher class, creating a descending list. The effect of such an allocation structure is to never refuse requests for the highest fare product, as long as the flight capacity ceiling has not been reached. In contrast, *partitioned* designs treat fare class booking limits as distinct terms. Because no further attention is paid to the class revenue ordering once seat allocations have been made, atypical demand arrival patterns might "close" a higher yield class (bring its availability to

[7]Some formulations replace the dependent variable, BTC, with TBD, or total bookings at departure.
[8]Roughly, the number of seats withheld from use by all lower fare classes. The mathematical language of Section 2.2.2.1 clarifies the notions of booking limits and protection levels.

zero) before all those below it have been filled. Table 2.1 illustrates how the same protection levels would be represented under nested and partitioned hierarchies. Partitioned structures, while adopted in many theoretical studies, can produce surprisingly poor results, at times underperforming an unmanaged system. As only nested architectures will be examined in the simulation study, partitioned controls receive no further attention.

Fare Class	Protection Levels	Nested Booking Limits	Partitioned Booking Limits
Y	10	100*	10
B	15	90	15
M	30	75	30
Q	-	45	45

*** Authorized capacity = 100**

Table 2.1. Nested Versus Partitioned Booking Limits Example

2.2.2.1 Leg-Based Approaches

The class of Expected Marginal Seat Revenue (EMSR) heuristics, introduced by Belobaba (1987), takes a probabilistic view of the inventory management problem. Given two fairly nonrestrictive statistical assumptions on the underlying booking process,[9] EMSR analysis begins with an assessment of the likelihood, $\overline{P}$, that realized demand for a fare class will meet or exceed a trial inventory allocated to the class:

$$\overline{P}_i(S_i) = \Pr\{X_i \geq S_i\} \tag{2.2}$$

where

X_i is the demand for fare class i;[10]

S_i is the number of seats open for reservations in class i.

As noted above, several forecasting techniques are used to derive demand density curves from empirical observations. Typically, a quasi-normal distribution emerges,

[9]Namely: a) demand independent across fare classes; b) demand uncorrelated over booking periods. See Chapter 6 of Belobaba (1987) for discussion on the risk and consequences of violation.
[10]In this and all later examples, fare classes are referred to by number (with the highest yield class = 1) rather than reservation letter codes, so the hierarchy is immediately apparent.

perhaps truncated over extreme values. But the right-hand side of (2.2), seen to be the area to the right of S in the cumulative distribution function associated with the random variable X, does not assume normality, nor any other particular density shape. Naturally, $\overline{P}$ is a decreasing function of S—as the assigned fare class inventory is increased, the chance that booking requests will surpass the limit falls.

When there are only two fare classes, seat allocation reduces to deciding on the number of higher fare spaces to protect from discount sales. Conforming to the ubiquitous optimization principle of microeconomics, total revenue will be maximized only if each revenue source—the two fare classes, here—contribute equally on the margin.[11] For were one, say class 1, expected to produce a greater marginal revenue given one additional seat, it would be advantageous to shift the current division and devote a larger share of resources to the upper class. The expected marginal revenue, *EMSR*, of making the Sth seat available to class i is:

$$EMSR_i(S_i) = F_i \cdot \overline{P}_i(S_i) \tag{2.3}$$

with F_i the fare in class i.

With a nested inventory scheme, the optimal *protection level*, S_2^1, can be found as the largest integer which meets the rule:

$$EMSR_1(S_2^1) \geq F_2 \tag{2.4}$$

The *booking limit* for class 2 is then set as the residual, unprotected capacity.

For fare class hierarchies with three or more levels, two versions of the basic EMSR method are available (Belobaba, 1992a). EMSRa is a straightforward extension of the two class case. With k nested classes, (2.4) expands to the set of conditions:

[11]Or as nearly as possible, given the discrete setting.

$$EMSR_i(S_j^i) = F_j \tag{2.5}$$

$$\forall i < j; i = 1,\ldots,k-1;\ j = 2,\ldots,k$$

where S_j^i is the number of seats protected *for* class *i*, *from* class *j*.

By subtracting the individual protection levels for each lower class from the authorized capacity, *C*, the booking limit, *BL*, for class *j* is:

$$BL_j = C - \sum_{i<j} S_j^i \tag{2.6}$$

$$j = 2,\ldots,k$$

From (2.5), EMSRa arrives at total protection levels for classes 1,...,*k*-1 by finding a *partial* protection level for each lower class in turn. In contrast, EMSRb considers the subsets of consecutive fare classes $(1;1,2;1,2,3;\ldots;1,2,\ldots,k)$ [12] and, in one step, decides how many seats will be *jointly* protected for the subset, from use by all lower classes. This joint protection level is denoted by π_n, for the subset whose last element is *n*. Within a subset, demand forecast data are collapsed to a single set of distribution parameters and expected marginal revenue calculations are based on an average fare, constructed using the mean demands as weights:

[12] Although the two EMSR varieties calculate class 1 terms, S_2^1 and BL_2, in the same manner (refer to [2.5], [2.6]), the class subset, $\{1\}$, is also included in the general EMSRb equations which follow.

$$\overline{X}_{1,n} = \sum_{i=1}^{n} \overline{X}_i \tag{2.7}$$

$$\hat{\sigma}_{1,n} = \sqrt{\sum_{i=1}^{n} \hat{\sigma}_i^2} \tag{2.8}$$

$$F_{1,n} = \frac{\sum_{i=1}^{n} F_i \overline{X}_i}{\overline{X}_{1,n}} \tag{2.9}$$

$$n = 1,\ldots,k-1$$

where

$\overline{X}_i$ is the mean forecast demand for class i;

$\hat{\sigma}_i$ is the sample standard deviation for class i demand;

F_i is the class i fare;

$\cdot_{1,n}$ is the corresponding aggregate term for the subset whose last element is class n.

Protection levels are set so that:

$$F_{1,n} \cdot \overline{P}_{1,n}(\pi_n) = EMSR_{1,n}(\pi_n) = F_{n+1} \tag{2.10}$$

$$n = 1,\ldots,k-1$$

Booking limits then follow directly, with:

$$BL_n = C - \pi_{n-1} \tag{2.11}$$

$$n = 2,\ldots,k$$

EMSRb improves on EMSRa by greatly reducing execution time, while producing similar revenue results. The ranges in (2.5) and (2.10) reveal that EMSRa calculations vary with the square of the number of fare classes. EMSRb complexity is linear in k. With ten fare classes, the work ratio is 1:5 in favor of EMSRb. Over an entire flight network, this translates to a significant savings.

In closing, it should be remembered that the EMSR models are heuristics, albeit very good ones. Curry's (1990) Optimum Booking Limits (OBL) approach finds exact protection levels by calculating the joint revenue function for each fare class subset, which EMSRb approximates by averaging the individual class parameters with (2.7)-(2.9). This task requires computationally intensive numerical methods including multidimensional convolution integrals. It has been argued that the small divergences between EMSR and OBL which would be observed in actual RM setting do not justify the burden imposed by the theoretically superior OBL (Belobaba, 1992a).

2.2.2.2 Virtual Nesting and Network Optimization

Although the limitations of leg-based inventory control are underscored in the large hub route systems established by most major carriers, the motivation for a new approach to RM can be seen with just a three city network. By representing revenue potential for each class on a flight leg with a single dollar value, and maximizing the revenue for each leg in the network independently, leg-based control produces suboptimal system revenue results whenever a single leg serves more than one OD market. A short example will illustrate this assertion.[13]

Imagine that the entire network for a small carrier consists of three cities: Boston (BOS), Chicago (CHI), and Seattle (SEA). Furthermore, range limitations on the carrier's fleet preclude transcontinental nonstop service. All flight legs are then in one of two sectors: BOS-CHI or SEA-CHI. It will be helpful to refer to the two shorter-haul OD markets joined by these legs as *local markets*. The remaining coast-to-coast market, BOS-SEA, is served by paths which combine one leg from each sector. The carrier coordinates its departures to allow the construction of such paths and BOS-SEA can be called a *connecting market*. As a result, each flight leg may have both local market and connecting market passengers.

Two fare products are offered in every market: a full-fare Y class ticket, and a discount B class ticket requiring a 7 day advance purchase. Prices differ by market, though, with a Y ticket costing $900 in BOS-SEA and $500 in both local markets. Similarly, the B prices are $600 and $200 for the connecting and local markets, respectively.

[13]More detailed examples may be found in Smith (1992) and Belobaba (1994).

Application of the EMSR formulas requires single revenue values to be assigned to Y and B classes. Suppose these are set at the average fares for the two markets which use each leg, or $700 for Y class and $400 for B class. Protection levels and booking limits are then determined by combining these values with forecast demands by fare class and comparing expected marginal seat revenues. Passengers traveling in either the connecting or local market place requests for a fare class on the leg and are accepted or denied subject to these control guidelines.

But, by basing reservation controls strictly on fare class, the carrier ignores the actual contributions of the two passenger markets. In these pricing conditions, a full-fare local passenger is worth less in revenue terms than a discount connecting passenger, but the entire leg seating inventory is open to the former (in a nested control structure) while the latter must contend with protection levels. This may produce the undesirable result of refusing a $600 B class connecting passenger, because the discount booking limit has been reached, but continuing to sell $500 full-fare Y class local tickets. While the reservation system infrastructure benefits from a common fare class hierarchy, direct application of this hierarchy to the purposes of revenue optimization may be inappropriate.

Intuitively, a better control method would refer to the actual revenues associated with each fare product served by the leg, across all OD markets. As early as 1983, major carriers have experimented with such a fare product-based control policy, through a procedure known as *virtual nesting* (Smith, 1992). Due to the enormous number of fare class / market combinations offered by even a small airline, virtual nesting identifies groups of similarly-valued fare products, rather than attempting to control each fare product individually. In the above example, the prices for the Y-local and B-connecting products might be considered sufficiently close that they could be combined into one *bucket*, the control unit in virtual nesting. Each of the Y-connecting and B-local products merit a separate bucket.

The "virtual" in virtual nesting stems from the continuing use of the traditional fare class codes in the reservations interface visible to external distribution channels such as travel agents. Only within the carrier's internal reservations computer system does a virtual link exist which indexes the fare class for a given market to the global hierarchy of buckets. When a request for a fare class arrives, availability is determined by mapping the request to a bucket which controls the leg inventory allocation for all fare products of roughly the same revenue value.

Virtual nesting seeks to account for revenue variations within fare classes to promote the RM mission of system-wide revenue maximization. At its heart though, virtual nesting still employs leg-based control, albeit "manually" enhanced to refocus the optimization effort on the revenue objective. In contrast, *network* or *OD-based* control methods propose to simultaneously manage all flight leg inventories given the flows of OD market traffic over an entire system network. The variety of formulations advanced in the last few years, and their comparative strengths and weaknesses in an operational setting, cannot be adequately addressed within the scope of this thesis. Again, Williamson (1992) offers a comprehensive survey of the recent research in this area. Here, only the basic rationale for OD-based control will be presented, followed by the common criticism leveled against this class of control methods.

In the example above, it was suggested that simple (non-virtually nested) leg-based control may prescribe the refusal of a discount connecting passenger to preserve space for a full-fare local passenger who may in fact yield a smaller total revenue. Implementation of virtual nesting would then correct this perverse outcome. However, even faithful adherence to virtual nesting often produces suboptimal system revenues.

Although virtual nesting explicitly incorporates revenue contribution, in some sense, the "cost" of supplying each fare product service is not captured in the system of buckets. Specifically, the system seat resource cost for connecting and local passengers are implicitly treated as equivalent, since from a leg-based perspective, a passenger from any market occupies one seat on the leg under consideration at any given point in time. However, stepping back to the network level, the connecting passenger must also be accommodated on the other legs of the itinerary.

This *network effect* leads to problems in the valuation of true system revenue contribution under certain demand conditions. If acceptance of a connecting passenger causes the subsequent refusal of a "down line" local passenger on the next flight leg, it seems misleading to credit as incremental revenue the entire fare paid by the connecting passenger, ignoring the revenue lost in the local market. Generically known as "opportunity cost" in microeconomics, this foregone revenue is referred to as *displacement cost* in the airline RM context. Network control bypasses the displacement cost issue by jointly optimizing inventory allocation for each fare class / OD market combination in all flight legs of the network.[14] The recognition of displacement cost becomes particularly important when two local passengers collectively provide more revenue than a single connecting passenger in the same fare class.

While its purpose is clear, the fundamental premise of network control—that demand can be forecast reliably on a fare class / OD market basis—is open to debate. From a technical standpoint, these numbers are very small and may have forecasting errors which are a significant fraction of their average values. Perhaps more importantly, the OD construct itself may be flawed on a conceptual level. Conceivably, modelers are attempting to disaggregate demand information to an inappropriate scale—that is, beyond the level at which passengers actually make their travel decisions. Rather than improving the accuracy of forecasts and subsequent protection levels, OD decomposition would then confuse RM analysis and contribute only a false sense of precision. At this point, while the concerns of those endorsing network optimization are well taken, it remains to be seen if such routines will indeed provide significant revenue benefits under actual demand conditions, given the current vague understanding of the underlying passenger behavior. This thesis establishes a simulation environment capable of exploring these questions.

2.3 Air Travel Demand Modeling

This section offers a brief introduction to the application of disaggregate demand models to air travel. First, a review of the fundamental concepts of discrete choice theory is presented. This background leads into the discussion of a specific logit model formulation designed for the air mode. The logit proposal is then contrasted with the PODS passenger choice modeling approach, which is described in Chapter 3.

Discrete choice theory extends the utility concept of microeconomics to choice situations in which the set of alternatives is not continuous. A common utility function for all alternatives gives the utility value of each alternative on an ordinal scale. Loosely speaking, the utility function returns, for each alternative, a weighted measure of all attributes relevant to the decision. One common application of discrete choice models considers the choice of travel mode for commuting trips. If the only attributes which

[14]Heuristics also exist for the inclusion of displacement cost in leg-based control methods (see Tan [1994] and Belobaba [1994]).

affected the choice of mode were travel time (*TT*) and travel cost (*TC*), with *TC* the more important attribute, a set of basic utility functions, *U()*, might look like:

$$U(car) = TT(car) + 2TC(car)$$
$$U(bus) = TT(bus) + 2TC(bus)$$

The alternative with the higher utility value would be preferred by a traveler with the valuations of *TT* and *TC* expressed in the utility functions.

However, in practice, not all of the influential attributes can be observed by a modeler of the choice process. As a result, under the random utility approach of discrete choice theory,[15] a stochastic term is added to the existing deterministic component of the utility function. The specification of the stochastic term determines the class of choice model. The logit choice model is used in many modeling contexts and assumes that the difference of the stochastic terms for two alternatives is logistically distributed.

Hansen (1990) advances a logit choice formulation in his model for passenger choice of airline in domestic markets. The deterministic portions of the additive utility specifications, V, for direct (d) and hubbing (h) carriers, respectively, are:[16]

$$V^d = \alpha F + \varphi_0 \ln(v) + V_{dpref} \tag{2.12}$$

$$V^h = \alpha F + \sigma CIRC + \varphi_1 \ln(v_{max}) + \varphi_2 \ln(v_{min}) \tag{2.13}$$

where
F is the average fare;
v is the direct frequency;
v_{max} and v_{min} are hub frequency terms (allowing for different service levels on each leg);
$CIRC$ is a circuity factor;
V_{dpref} is a choice-specific constant, measuring passenger preference for direct service, all else equal.

[15]Ben-Akiva (1985), p. 55.
[16]Hansen (1990), p. 31.

Service frequency, evaluated separately for the direct and hub carriers, stands as the lone strategic variable in Hansen's model. Fares are held constant under a "price-taker assumption." Frequency appears in logarithm form to reflect the notion of frequency saturation—at high levels of frequency, a marginal flight contributes relatively little to the overall utility of the carrier's service.

The estimated carrier market share model of this (and any other) logit choice model is attractive in its computational simplicity, with the fraction of passengers selecting each available service given as:[17]

$$S_{ij}^{m} = \frac{\exp(V_{ij}^{m})}{\sum_{n=1}^{C} \exp(V_{ij}^{n})} \tag{2.14}$$

where

S_{ij}^{m} is the market share of carrier m in market ij;

V_{ij}^{m} is the deterministic portion of the utility function for travel on carrier m in market ij;
C is the total number of carriers in the market.

Hansen's ultimate goal is to simulate the performance of a reduced version of the domestic air transportation system. The discrete choice framework developed here is only one ingredient in the carrier profit function, which also include cost and average fare components. Since the concern here is only with the discrete choice model, the rest of his research will not be analyzed.

This logit formulation has been presented as representative of previous efforts in the discrete choice modeling of air transportation services. Indeed, Alamdari (1992) conducts a review of twelve recent independent studies, devoted to the modeling of airline and airport choice, which share the same general logit form with different utility specifications. There are three readily identifiable benefits of adopting the logit approach for these purposes.

First, as already noted, the resulting market share equation has a closed form solution. Most other specifications for the stochastic term produce intractable expressions

[17]Hansen (1990), p. 31.

for market share. Also, developing an *a priori* behavioral model which underlies a utility specification is more responsible than the alternative of culling data to find a combination of explanatory variables which happens to produce an acceptable explained sum of squares in conventional regression. Finally, logit appears to be very flexible—the range of expressions which might be inserted in Equations 2.12-2.13 is limitless. Given these advantages, the popularity of logit models hardly seems surprising.

However, all traditional models of discrete choice, including logit, handle one decision attribute very poorly—time. Though less important for other transportation modes characterized by regular headways, the discrete timing of aircraft departures is absolutely crucial in modeling carrier choice and, to an even greater extent, flight choice within carrier. In many market settings, absolute frequency counts may be only very rough proxies for service quality, and in turn, passenger utility.

It may seem that this failing could be corrected by merely adding enough indicator variables to the utility specification. For instance, if studies of passenger preference reveal that placement of a departure at the 5:00 PM peak should receive a special utility weight, this could be incorporated by an alternative-specific binary variable which takes a value of 1 if the flight leaves between 4:45 PM and 5:15 PM, and 0 otherwise. Capturing the entire preference map over the operating day with such variables would be a time-consuming exercise, but seems theoretically possible.

Yet the problem has deeper roots. For a carrier to achieve the full utility bonus in the eyes of the consumer, it is not enough just to have a 5:00 PM departure. The time slot must also not be served by competing flights. In short, the relative timings decide the utility rankings. The existence of these interaction effects renders impossible a logit-based solution to the modeling of flight choice. An entirely new choice framework must be introduced to handle the temporal dimension. Boeing's Passenger Origin / Destination Simulator (PODS) offers such a framework.

*   *   *

This chapter has first placed RM within the broader context of air transportation supply activities. The defining elements of the airline RM problem are then identified. An introduction to the forecasting process is included before reviewing the optimization methods available to set the capacities allocated to each fare class. Three optimization scopes are described and a detailed analysis of leg-based algorithms precedes a conceptual discussion of virtual nesting and network-control approaches. The final section considers discrete choice modeling of air travel demand and stresses a shortcoming of traditional logit formulations—the treatment of departure time. The PODS model of Chapter 3 explicitly captures this temporal dimension.

3 The Passenger Origin / Destination Simulator

Modelers at the Boeing Commercial Airplane Group have had a long-standing interest in the behavioral processes which govern passenger demand for air travel. By the early 1990s, Boeing had produced an integrated probabilistic model, known as the "Decision Window" (DW) system,[18] which incorporates various research findings on how passengers choose one airline service option over another, given a set of attributes associated with each competitor in the market. The DW approach makes its primary contribution to the field of airline demand modeling by explicitly recognizing the importance of the scheduling dimension. In contrast to earlier logit-based efforts, relying solely on absolute frequency counts, DW contends that the relative timings of flights contain significant explanatory power for market share analysis.

In implementing this hypothesis, DW first defines the preferred travel period for each individual passenger, bounded by the earliest departure time and latest arrival time that suit the passenger's plans. This construct, the "window" which gives the model its name, is then checked against the scheduled travel times available in the market and only those alternatives lying entirely within the passenger's window will be considered in later stages of the decision process. The finer distinctions, concerning path quality (i.e., nonstop versus connecting), airline image, and even aircraft preference (i.e., propeller, narrowbody jet, widebody jet), capture important secondary factors and allow for a more accurate representation of actual competitive situations.

DW was originally intended as a high-level fleet planning and scheduling tool. Using DW, modelers were able to experiment with alternative departure placements and aircraft types to determine a flight schedule which would deliver an acceptable level of service, as measured by average load factors. However, despite the quantum advance in DW's treatment of competing schedules, the macroscopic view of the model leaves it unsuitable for closer examination of the relationship between unconstrained path preference and the realized passenger flow. The omission of prices particularly limits the range of planning scenarios which may be tested with DW. To expand the applications of the DW approach, Boeing has created the Passenger Origin / Destination Simulator (PODS) which incorporates prices and multiple fare classes within the broader passenger assignment framework pioneered in DW. Currently, Boeing and MIT's Flight

[18]Boeing Commercial Airplane Group (1993).

Transportation Laboratory are embarked on a collaborative study to explore the modeling of RM functions in the PODS environment. As the first MIT publication to evolve from this work, this thesis takes as its purposes the introduction of the PODS methodology and a review of results from small-scale modeling experiments employing basic RM principles.

The remainder of this chapter is devoted to the first objective. Following an outline of PODS' macro-scale structure, each of the three main modules—the stochastic demand generator, passenger assignment, and RM submodels—are described in turn.

3.1 High-Level Architecture and Assumptions

Figure 3.1 illustrates the organization of the component models in PODS. The first and third boxes comprise the routines which produce both aggregate market demand and its distribution over the booking horizon. Box two includes the forecasting and inventory management functions of the RM system. The fourth contains the passenger assignment logic as well as the necessary interaction with reservation control as fare product availabilities are adjusted (decremented, and perhaps closed) after each successful assignment. The final box in the diagram outputs the operating results of the simulated departures, in terms of revenue and traffic categorized by fare class and passenger segment.

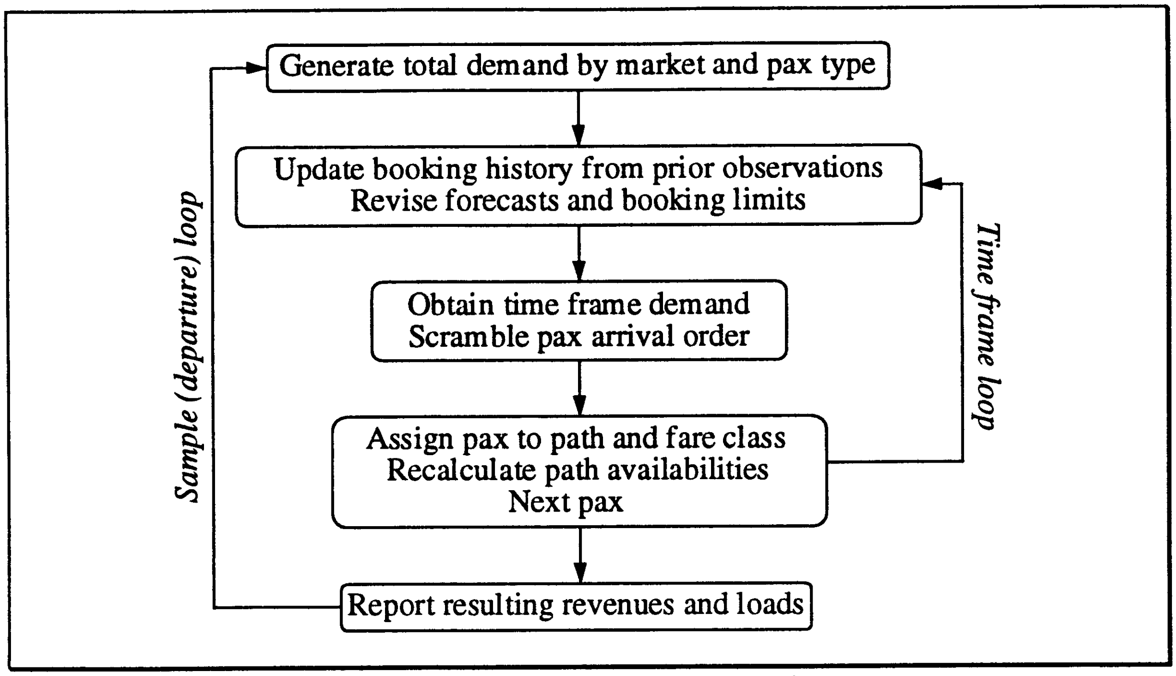

Figure 3.1. PODS System Flow[19]

Two loop structures link the five elements. The outer loop iterates on day of departure, known as a *sample* or *observation* in PODS terminology. Each day, all scheduled paths receive bookings simultaneously so the actual counter variable is a set of flight departures. For instance, a single-market experiment with 100 observations and four available paths would simulate 400 individual departures. Nested within each observation, the time frames of Section 2.2.2 are represented with a second loop which contains calls to the RM procedures, effecting dynamic revision.

Several types of input are necessary for model execution. *System-level* inputs affect the entire simulation and include the items classified in Figure 3.2. *Airline-specific* features apply in all markets containing legs operated by the airline and are listed in Figure 3.3. *Market, path,* and *leg* information, in descending order of detail, complete the specification of the network (Figures 3.4 to 3.6).

[19]Figure adapted from Hopperstad (1995), p. 5.

- sizing variables; total number of:
  - airlines
  - legs (path segments)
  - markets

- behavioral coefficients: costs assigned by the segmented passenger types to flight option attributes

- fare and reservation structure descriptors:
  - number of:
    - fare classes
    - restriction categories
    - observations stored in forecaster
  - restriction definitions:
    - last booking period available
    - flags indicating which restrictions apply to each fare class
  - specification of booking tables (curves)

- stochastic factors

- experimental design inputs:
  - number of:
    - total samples
    - samples deleted prior to collection of summary statistics
  - random seed flag (allows replication of demand conditions)

Figure 3.2. System Inputs

- RM system definition:
  - optimization algorithm
  - forecasting method
  - fare allocation scheme for paths containing multiple legs
  - options affecting estimation of forecasting parameters

- coefficients of carrier preference (airline image) by market

Figure 3.3. Airline Inputs

- choice between Boeing market research and user-calibrated surveys on:
  - mean schedule tolerance (flexibility with respect to elapsed flight time)
  - time-of-day demand profile
- competitive environment:
  - number of paths available in market
  - mean market demand by passenger type
  - fares

Figure 3.4. Market Inputs

- scheduled departure and arrival times
- path quality index
- component leg identifiers

Figure 3.5. Path Inputs

- booking (airplane) capacity
- distance
- initialization for RM computations:
  - forecast mean demand
  - forecasting error
  - attributed fare

Figure 3.6. Leg Inputs

As will be demonstrated in the subsequent subsections, the flexibility of the PODS framework allows experimenters to represent most significant competitive dimensions in complex market settings. Nevertheless, certain fundamental assumptions of the broad modeling environment, which simplify aspects of booking and passenger choice processes observed in practice, should be noted here. First, PODS assumes a *stationary process*, in that the parameters describing booking behavior are taken as uniform over the simulation period. This implies that seasonal trends, an important confounding factor in actual airline forecasting, are not incorporated. Furthermore, given the day-to-day traffic variability typically observed within a week, PODS should be thought of as simulating the performance of a flight on a single day-of-week. Sample size, then, specifies the number of weeks over which the departure will be modeled.

Another restriction which limits direct application of PODS to the airlines is the omission of reservation phenomena including cancellations prior to the day of departure, no-shows, and overbooking.[20] Due to their importance in the forecasting and optimization routines of modern RM systems, Boeing and MIT, with carrier input, are currently evaluating appropriate models to manage these complications, due to be implemented in the near future. Such improvements would be expected to refine the analysis contained here by clarifying the benefit achievable by RM when confronted with various levels of cancellation activity and no-show rates.

Finally, a methodological issue arises as it is noticed that PODS places its two primary processes, passenger assignment and RM availability control, on the same time scale. Intuitively, it might be argued that passengers do not make decisions about airline service at each departure opportunity. For instance, there may be a kind of inertia influencing a business traveler who takes frequent trips in the same market. Once an initial flight selection is made, perhaps to conform with the schedule of the first meeting, successive visits may well be partially determined by historical experience and familiarity with options previously chosen. If true, modeling of the passenger decision process as an independent assessment of the entire schedule before every reservation would be inaccurate. On the other hand, tactical inventory management does indeed necessarily occur on a day-to-day, sub-departure basis. Despite the apparent conflict, this modeling concern is less significant than the two described above. First, given the current limited understanding of the true passenger decision process, a more suitable alternative for integrating a conventional demand model structure with RM controls is not readily at hand. Second, the PODS approach would be more appropriate if the prospective passenger population were a large pool of infrequent travelers. This condition may be met to varying degrees in some leisure markets.

---

[20]Or, more precisely, "oversales." Because of the possibility of cancellation and no-shows, overbooking will not always be problematic. Only when a flight is oversold—that is, when the number of passengers appearing at the gate exceeds the capacity of the available aircraft—must the carrier resort to denied boarding procedures.

While these reservations should not be lightly dismissed, the basic strength of the modeling framework is apparent. All previous booking simulation efforts have taken as given the demand (distribution) which will seek accommodation on a certain path, under a set of alternative inventory control methods. In contrast, PODS endogenously determines carrier and individual flight demand by means of a passenger choice process which captures competitive factors particular to air travel, including the very treatment of this demand as prescribed by the industry's RM optimization principles (Figure 3.7). This cycle can be seen as a variant of the traditional sequential trip generation-distribution-route selection model of ground transportation demand analysis context. Through the RM mechanism, an additional feedback effect is introduced which causes behavioral characteristics and choices of current passengers to affect the capacity airlines are willing to supply in the future, given the historical nature of the forecasting exercise.

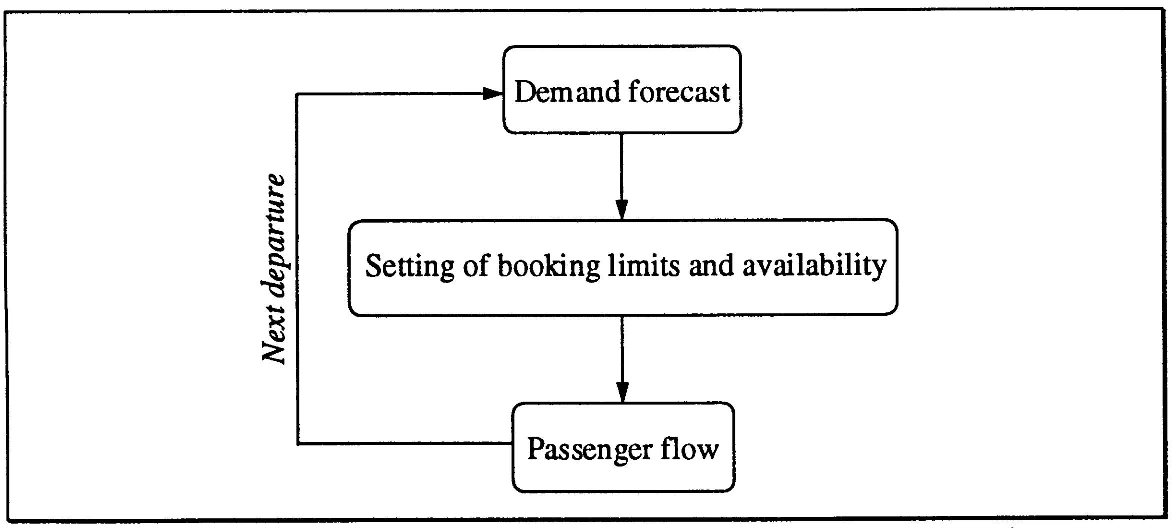

Figure 3.7. Forecasting-Revenue Management-Flight Choice Cycle[21]

3.2 Demand Generation

Figure 3.8 summarizes the PODS demand generation process. The first step is the calculation of the number of potential passengers, by segment, who will attempt to travel in an OD market. As might be expected, the starting point for this procedure is the initial value of the forecast market demand contained in the scenario input file. To incorporate random deviations around the average, PODS follows the common industry practice of assessing a variability measure which depends on the magnitude of the mean. Two alternative forms for this stochastic term have been advanced, referred to as *k*- and *z*-

[21]Hopperstad (1995), p. 2.

factors, and after a brief definition of each, it will be seen that PODS employs both sorts of factors in a two-phase model. The primary phase (first three boxes of Figure 3.8) determines the aggregate "propensity to travel" in the market and the day-to-day variability around this average demand. As a secondary objective (final two boxes), correlations among individual reservation requests are captured.

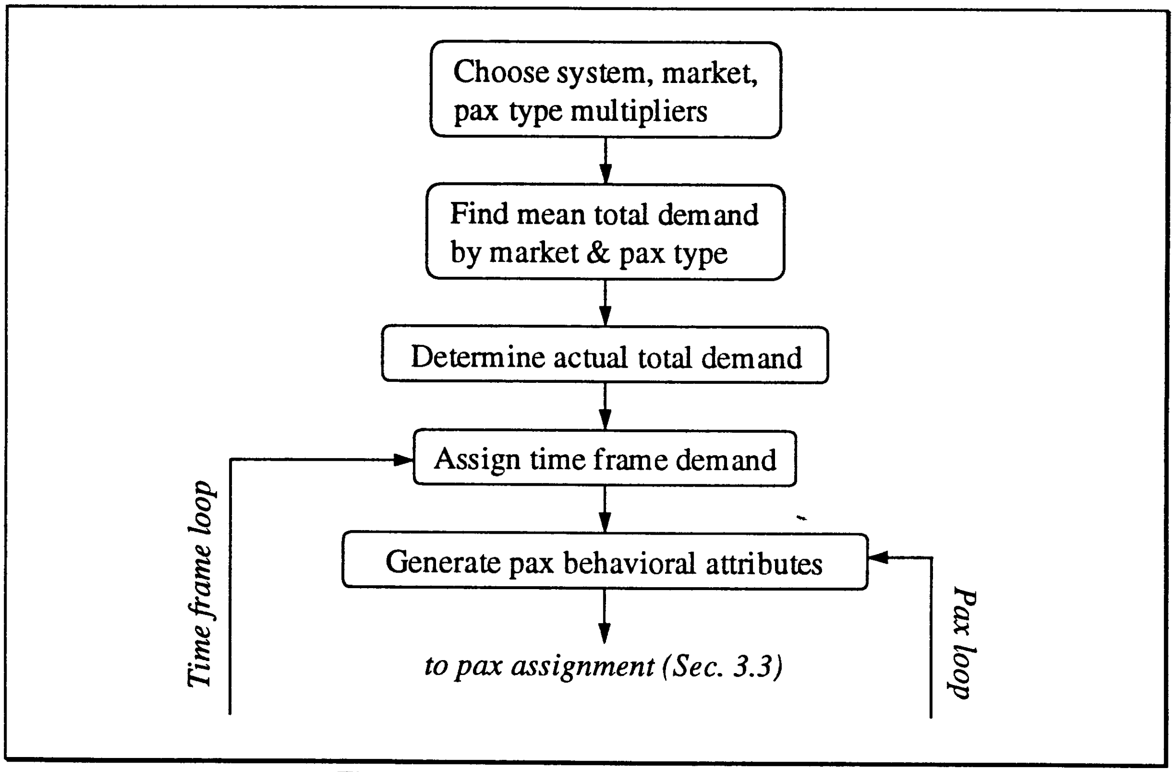

Figure 3.8. Demand Generation Flow[22]

Application of a k-factor supposes that the standard deviation, σ, of a random variable can be found as a constant times the mean, μ, or $\sigma = k\mu$. From empirical analysis of airline demand data, a k value close to 0.3 is typically derived, subject to special market conditions. If demand is modeled as approximately[23] normally distributed, knowledge of σ quickly gives the range ($\mu \pm \sigma$) in which $\frac{2}{3}$ of the observations would be expected to lie—for an average of 100, the bounds for this range are 70 and 130. It seems plausible that random variables with high means generally exhibit larger absolute discrepancies over multiple observations. However, the z-factor approach offers another possibility for the relationship by specifying that it is the variance, σ^2, which is directly proportional to the

[22]Hopperstad (1995), p. 11.

[23]In the case of random variables for which negative values would be nonsensical, the density function must be truncated at zero.

mean, or $\sigma^2 = z\mu$. Recent Boeing studies[24] suggest this form, which implies a smaller difference in the random variability for processes with low and high mean values, may be more stable for the modeling of air travel demand.

PODS builds a system of nested k-factors, from three user-input k-factor constants, to generate the *mean total demand* for passenger type *t* in OD market *m*, d^{μ}_{mt}, for a given observation. A system k-factor, *skf*, applies to all passengers in the network and can be thought of as reflecting the high travel periods universally observed near major holidays, as well as systemwide traffic slumps, as may be seen in late winter. A market k-factor, *mkf*, quantifies the extent to which unusual demand activity in individual passenger type move together. The values for the first two common k-factor inputs, then, describe the strength of the positive correlation among markets within the system, and among passenger types within a market. The residual within-type variability is represented by a third k-factor, *tkf*. To derive the k-factor for a specific market / passenger type, each of these is multiplied by independently drawn standard normal variates, denoted as *z*.[25] The terms are then combined to produce a single multiplier, $kmult_{mt}$ (Equation 3.1). After verifying that this grand multiplier is nonnegative, it is applied to the initial forecast, d^{1}_{mt}, to obtain d^{μ}_{mt} (Equation 3.2).

$$kmult_{mt} = 1 + (z \cdot skf) + (z \cdot mkf) + (z \cdot tkf) \tag{3.1}$$

$$d^{\mu}_{mt} = kmult_{mt} \cdot d^{1}_{mt} \tag{3.2}$$

In order to select the *actual total demand* for the sample, d_{mt}, a measure of the variability around d^{μ}_{mt}, is calculated using the first of two input z-factors, zf_1. In a similar manner to the steps followed above,[26]

$$d_{mt} = d^{\mu}_{mt} + \left(z \cdot \sqrt{d^{\mu}_{mt} \cdot zf_1}\right) \tag{3.3}$$

[24]Hopperstad (1995), p. 12.
[25]Not to be confused with a z-factor!

[26]This result must be nonnegative and fractional values are rounded to the nearest integer.

Given this total demand, which will appear at some point in the booking process, the allocation among time frames by passenger type remains to be determined. For the earliest time frame (furthest away from departure), the average demand is simply the product of d_{mt} and the corresponding incremental booking probability from the appropriate curve. The variance expression is derived by treating time frame assignment as a sequence of Bernoulli trials, where the probabilities of: a) booking in the current time frame; and b) waiting until a later period, are translated as success and failure, respectively. After randomization with the second z-factor,[27] the realized demand for the time frame is decremented from d_{mt}.

The remainder of the demand generation process describes the behavioral characteristics of those passengers arriving in the time frame. First, the members of each passenger type are combined to form one ordered list. For the first passenger on the list, four sets of attributes, related to price and perceived cost, are then determined:

- acceptable price ratio
- inferior airline cost ratio
- path quality index unit cost
- ticket restriction costs

The first two items are ratios applied to the *base fare*, typically taken as the lowest fare on offer in the market. However, if the lowest fare has arisen from a special sale and is not deemed indicative of the prevailing discount price with respect to which passengers make their travel decisions, the base fare may be entered separately. The acceptable price ratio marks the upper limit for the fare dollar price—a more expensive product will be dismissed out of hand, regardless of the accompanying restrictions and service quality. The other three give the attributed costs which are summed and compared in the passenger assignment phase (Section 3.3). All of the four parameters are calculated from averages input for each passenger type. Application of a generic attributed cost k-factor produces individual preferences, normally distributed around the mean values.

---

[27]The essential effect of the two z-factors is to determine the correlation in bookings across time frames—the smaller the ratio of zf_1 to zf_2, the lower the correlation, and no temporal correlation exists when $zf_1 = zf_2$. To preserve the integrity of the simulation, this correlation is not made available to the forecasting routines—it must be divined from the emerging booking history. A behavioral interpretation of these secondary processes can be found in Hopperstad (1995), pp. 14-15.

Subsequent passenger behavioral parameters are created similarly, but with one additional step. To model situations where groups of various sizes (families, business associates) seek accommodation, PODS occasionally allows a passenger to share the predecessor's cost sensitivities. This occurs with some small probability determined by the z-factor previously used in time frame distribution. However, the group concept is defined by behavior only—all passengers are still treated independently for booking purposes. That is, if a group of three find only two seats open in their preferred class, the group is willing to split, leaving the third to consider the other available options. To describe this *replanning process*, the general passenger assignment model is presented next.

3.3 Passenger Assignment

Within the passenger loop of Figure 3.8, once the individual behavioral profile has been generated, a service option is chosen following the procedure shown in Figure 3.9. As indicated above, PODS rates the acceptable price attribute as the overriding criterion in this decision—recognizing that the ticketed fare, as the most visible cost of the reservation process, has priority over all other ticket properties in the mind of the consumer. The randomized acceptable price ratio, when applied to the base market fare, determines the maximum out-of-pocket expense a passenger is willing to incur. Any flight for which the dollar price of the cheapest available fare product exceeds this limit is summarily discarded. Should this screen eliminate every flight, the passenger is *spilled* and (outside of the PODS setting) must either find another transportation mode or postpone the current plans for air travel. Spillage volume thus depends not only on the size of the imbalance between market demand and capacity, as would be expected, but also on the price sensitivity of prospective passengers, relative to the fare structure.

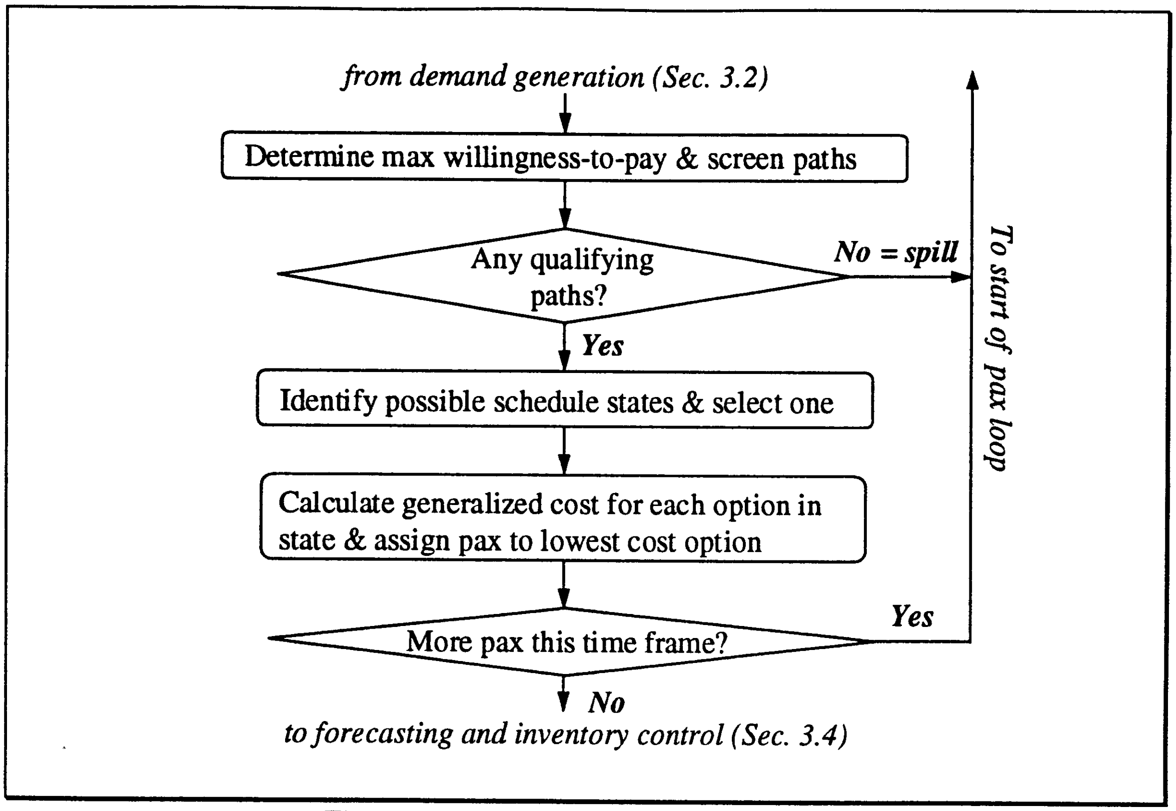

Figure 3.9. Passenger Assignment Flow[28]

More typically, one or more paths will have space available at an acceptable price, requiring the passenger to weigh the merits of each. Prior to this comparison stage, however, PODS makes its contribution in the area of temporal competition by introducing the mechanism of *schedule states*. A schedule state is a combination of paths which solves a passenger's trip planning problem, as introduced in Section 3.1. PODS represents each state using a string of binary indicators, with 1 and 0 indicating set membership and exclusion, respectively. Figure 3.10 illustrates the simplest nontrivial example of a flight schedule and its associated schedule states. With the two daily flights in the market as shown, a passenger may be willing and able to travel on Flight A only, Flight B only, or both. Note that the "empty set" state, which would describe the case in which none of the scheduled flights is feasible, has not been included. PODS implicitly assumes that all potential passengers in a market always find at least one suitable departure opportunity.

[28]Hopperstad (1995), p. 8.

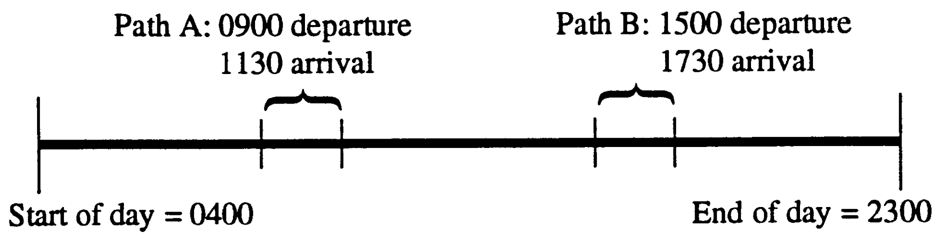

	Schedule States		Prior Probabilities	
State No.	A	B	Bus.	Leis.
1	1	0	0.48	0.47
2	0	1	0.03	0.08
3	1	1	0.49	0.45

Figure 3.10. Timeline and Schedule States List

The substance of the schedule state construct lies in the allocation of probabilities which govern the likelihood that the flight options for an individual passenger can be characterized by a certain state. In deriving these probabilities, PODS adopts the choice framework proposed by its forerunner, the DW model.

Both PODS and DW consult *time-of-day demand curves*, particular to each passenger type, to produce the distribution of decision window positions over the operational day. The standard curves plotted from Boeing market research exhibit well-defined morning and afternoon peaks—naturally more pronounced for the business segment—but the fall-off in traffic around these peaks is not as severe as that in similar profiles for ground transportation commuting trips. If desired, user-defined demand allocations, in half-hour increments, can be substituted to simulate special patterns.

Once a window's location is known, two additional elements are summed to determine its width. The first is a market-specific parameter which measures the lower bound on travel time, imposed by aircraft technology. This *market delta-t* (time) is typically computed as the duration of the shortest nonstop service offered by any participating carrier. It can be thought of as a sunk temporal cost in that all air passengers in the market confront this basic cost. Common to every flight option, the market delta-t is merely a scaling factor and irrelevant to comparative decision making.

The second component of window size is *schedule tolerance*, which represents passenger flexibility when considering the available flight times. Intolerant passengers, at the business end of the spectrum, accept little leeway around the market delta-t, resulting in a narrow window and a restricted set of possible paths. Leisure travelers, less concerned about the precise scheduling of arrival and departure, buffer the market delta-t with some level of slack time, generally leading to the inclusion of more options in their windows. Again, the default distributions for the tolerance terms can be overridden by providing mean inputs for each passenger segment.

The collection of finished windows is next conceptually overlaid on the flight schedule. PODS then assesses the *prior state probabilities* by counting the windows wholly containing the pertinent paths and dividing by the total number of windows in which at least one path lies (recalling the note above on the empty set condition). This procedure has two implications. First, due to the continuous form of the decision windows, any flight schedule comprising more than two paths includes certain impossible states. For instance, if Flights A, B, and C occur chronologically and have equal durations, a window containing A and C must also overlap with B, leaving state "101" with zero probability. Second, when only one path is available, considerations of time-of-day and schedule tolerance are moot as all market demand will be forced to seek space on the lone flight.

From this universe of states, PODS narrows the field of candidate states for an individual passenger to those containing at least one path with a fare product that passes the acceptable price criterion. The remaining states are reevaluated to obtain *conditional probabilities*, read as the chances that a combination of paths will fall within a randomly selected window, given the current passenger's price sensitivity.[29] In the example of Figure 3.10, if Path A were effectively closed to leisure traffic, the prior probabilities would be rescaled to show that States 2 and 3 now arise 15% and 85% of the time, respectively. Indexed against limits constructed from the probabilities, a number drawn from a uniform distribution over the range [0,1) then yields the state whose paths are the final contestants in the travel decision.[30]

---

[29]Window distribution and price sensitivity are likely to be correlated.

[30]This step is an artifact of Monte Carlo simulations such as PODS: clearly the passenger does not "roll a die" to decide his proper state! Probability here should be interpreted as relative frequency counts for a large sample. A passenger is simply more likely to choose from the options of a high-probability state.

Once the groundwork is laid during the decision window modeling phase, the final stage of the assignment process is straightforward. The perceived cost coefficients listed in Section 3.2, collectively constituting a utility function for air travel, are applied to all available fare products for paths in the chosen state, and the option with the lowest generalized cost receives a booking request. Should the passenger exhaust the paths in the state without finding an acceptable fare, replanning occurs by the selection of a new qualifying state. Spillage results only if no path, in any state, offers an available fare product meeting the maximum price restriction.

3.4 Forecasting and Inventory Control

The evolution from DW, a model which disregarded prices and the reservation process itself, has meant that all forecasting and inventory control within PODS takes place outside of the main program, which consists of the demand generation and passenger assignment routines. The simulation proper deals only in path availability and is indifferent to whether the information returned from RM function calls depends on leg- or path-based calculations. Such separability is an asset because it permits experimentation with any existing or theoretical RM system, or a customized method combining elements from various systems. Although the competitive analyses in this thesis cover only EMSR optimization and the pick-up forecasting approach, these are not "hard-wired" in PODS, and extensions into network-control algorithms are anticipated in the ongoing research plan.

PODS supplies two classes of inputs to the forecasting and inventory management routines. It will be recalled from the Chapter 2 review that the booking history, created here by the generation and assignment subsystems, stands at the foundation of all RM efforts. The presence of stochastic demand elements warrants careful statistical handling of the booking record, even under a stationary process. The other natural requirement is the fare product hierarchy, unchanging over the course of the simulation. Besides the price levels, which are directly incorporated in booking limit calculations, advance purchase restrictions both influence time frame demand patterns and act as additional availability constraints.

There are three components to the RM submodel:[31]

- booking parameter estimator
- availability calculator
  - forecaster
  - fare calculator
  - optimizer
- availability re-calculator

Because these components are called at different points in the time frame cycle, the RM submodel is not conducive to a flow chart representation. Instead, a few words on the function and input / output specifications for each routine follow.

The booking parameter estimator identifies the historical booking records of untruncated observations and attempts to reconstruct the booking curves which have generated this realized demand. To replicate the uncertainty of the real world RM process, PODS maintains a "Chinese wall" between the demand generation and booking parameter estimation routines. In other words, the demand inputs are unknown to the RM module—the booking trends must be divined from the emerging booking history, clouded by multiple stochastic elements (Section 3.2). For each complete untruncated observation stored in the reservations system, the booking parameter estimator calculates the cumulative fraction of total bookings received in each time frame. This statistic is found by dividing the bookings-in-hand at the end of the time frame by the total bookings received on the day of departure. The individual estimates for each observation are then aggregated by the averaging technique of exponential smoothing. Exponential smoothing assigns each observation a weight which is an exponential function of the observation's distance from the departure currently being forecast. At the end of this estimation process, the booking parameters are used to detruncate the other historical observations which closed prior to the day of departure (Section 2.2.2). For example, if the estimated booking parameter were 0.5 for the time frame in which a flight closed after receiving 100 bookings, the resulting detruncated demand would be 200.

The availability calculator, consulted only at the beginning of a time frame, first produces leg demand forecasts by fare class. Regardless of whether the regression or pick-up forecasting model is chosen, the process may include the entire stored history

---

[31]Hopperstad (1995), p. 17.

(detruncated and untruncated observations) or only the untruncated observations. In the second stage of the availability calculator, attributed leg fares for each fare class are determined by one of two methods. The first option sets the attributed fares equal to the fares for the nonstop market served by the leg. Alternatively, a distance-weighted average of the fares in all markets using the leg can be used to account for the differing fares among markets competing for the same leg inventory (naturally, the two methods generate the same results when the leg appears in a single market path). Under either approach, only one fare may be attributed for each fare class, as the control system remains leg-based. Finally, the leg demand forecasts and attributed fares are passed to the appropriate optimization routine. Nested implementations of both EMSR variants follow the procedures of Section 2.2.2.1. The output of the availability calculator is a set of fare class booking limits for the leg.

The availability re-calculator is called repeatedly during a time frame, after each booking request. Path availability is determined by checking the status of the requested fare class on each of the component legs. If any have been closed, flags indicating unavailability are set in all paths traversing the leg. If a request can be accepted, the availability re-calculator decrements the relevant fare class limits on each affected flight leg.

* * *

This chapter has first provided the motivation for the PODS modeling approach. Section 3.1 lists the classes of input to be defined by the user and illustrates the macro-scale architecture of the simulation. The next three sections describe the purpose and implementation of the main program submodels: demand generation, passenger assignment, and forecasting and inventory control. With this outline of the PODS model infrastructure, operational definitions for the simulation environment follow in Chapter 4.

4 The Operational Competitive Simulation Environment

The overview of the Boeing PODS model contained in Chapter 3 considered the interfaces between the three major program components, as well as the fundamental processes within each PODS component. For instance, the flow of information from demand generation to passenger assignment, and the underlying routines for stochastic demand generation itself, have been established. However, execution of PODS requires several additional user-selected inputs. These classes of inputs, and the values chosen for the simulation experiments in this thesis, are the subject of this chapter. The experimental simulation model which results from this further definition will be referred to as the Operational Competitive Simulation Environment (OCSE), to differentiate it from the generic Boeing PODS framework.

The first two sections describe the data requirements. Section 4.1 isolates the most pervasive and direct determinant of RM performance—the prevailing fare product hierarchy. Section 4.2 collects all other necessary inputs, explains the role of each in the simulation, and identifies the values selected for the base case simulation environment. To assess the effect alternative settings would have on later results, Section 4.3 conducts a brief sensitivity analysis for the critical noncompetitive parameter values.

4.1 Fare Product Structure

Chapter 2 discussed the dual nature of the RM effort. Specifically, revenue optimization relies on both the pricing of fare products and the allocation of capacity among these products. While the interrelationship between pricing and inventory control is an intuitive one in any field of microeconomic analysis, the complexity and dynamic nature of the simultaneous problem, in the airline context, often forces a sequential resolution of infinite iterations. That is, inventory on a future departure must be managed given that the prices for the current fare product assortment are established and cannot be incrementally adjusted to improve the optimization outcome. Fares do, of course, fluctuate considerably in the short term, in response to marketing initiatives from market rivals. But this reactive behavior does not allow for the joint setting of fares and inventory levels on individual departures with the goal of revenue maximization. Only on a strategic level, in the longer term, can a pricing policy be devised which is tailored to the characteristics of the market passenger base, as revealed by the previous inventory management history.

Validated then, in some sense, by industry practice, all simulation studies in this thesis assume a static fare environment.[32] Such an assumption actually implies relatively little loss of generality because path preference decisions and their consequences are better thought of as ordinal rather than cardinal processes. In other words, relative rather than absolute fares influence the behavior of passengers faced with multiple path and fare product combinations in reaching the desired destination. Furthermore, the robustness of the fare arrangement may be tested indirectly by varying the analogous passenger behavioral parameters. For instance, within the simulation, there are two ways to examine the effect of causing a high-yield fare product to be unaffordable for the average leisure passenger. The first is to raise the price of the fare product. The alternative would be to lower the average willingness-to-pay, or increase the price sensitivity, of the leisure passenger type. Both accomplish the same result but the second is easier to adopt in the PODS input framework. The results of the sensitivity analysis in Section 4.3, which explicitly considers variations in the behavioral parameters, can then also be interpreted as the effect of equivalent price changes. → No stimulation of demand by [illegible] fares

The reason for choosing the fare product structure as the unchanging base for sensitivity comparisons emerges in the strength of the sources consulted in its construction. After deciding on a relatively simple four class fare product line, a review of the on-line reservation service, EAASY SABRE, across a broad spectrum of markets, revealed a set of three restriction categories commonly used in the industry today: Saturday night stay, roundtrip ticketing, and refundability. The first two are binary indicators often found attached to discount fare products intended for purchase by leisure travelers. However, refundability may be modified by a percentage term—purchase price may be entirely (100%) refundable, 50% refundable, or nonrefundable (0%). This finer gradation distinguishes higher-yield options. Finally, typical advance purchase checkpoints of three, seven, and fourteen days were effective at the time of the fare product study (fall 1994).[33]

---

[32]References to RM innovation, then, concern the adoption of inventory control.
[33]Since then, many of the majors have raised the advance purchase fence of the cheapest excursion fare to 21 days. Under the booking curves used in the thesis, incorporation of the 21 day restriction would have a negligible impact on the results which follow.

Combining these ticket attributes in a hierarchy consistent with the majority of the surveyed markets led to the fare structure illustrated in Table 4.1. The spectrum runs from a full-fare unrestricted Y class product to a fourteen-day advance purchase, nonrefundable, excursion Q class fare. Differentiating the second-highest yield product in B class from Y class is a modest advance purchase requirement and a limited refundability provision. In the discount range, M class mirrors the restrictions of Q class but has a less severe advance purchase requirement of seven days. The reader will observe that the structure represented in Table 4.1 conforms with the convention that each member in a well-defined fare product list be superior to all lower-yield members, in every restriction element (i.e., excluding price). There exist no trade-offs in the restrictions alone—for example, an M class ticket is strictly preferred to a Q class ticket, because, in the only dimension in which they differ (advance purchase), the former has a less burdensome requirement.

Dollar prices for the fare products, also recorded in Table 4.1, were obtained similarly. However, the appropriate values were somewhat more difficult to identify from the EAASY SABRE quotes for two reasons. First, the fare environment in a few of the surveyed markets had clearly been affected by special competitive circumstances. At times, two fare products offered by the same carrier, seemingly identical in the restriction categories noted above, had quite different prices. Possible explanations for such anomalies involve other, unobserved restrictions, including limited duration sales or capacity controlled offers. After ascertaining the fare level more likely to prevail in more typical operations, there remained a significant spread in the absolute fares for equivalent fare products, across markets. This occurs because, even if all influential restrictions could be captured, many other market-specific attributes affect fares, such as length of haul, number and strength of competitors, and importance of the market in hub networks. To arrive at a single set of fares for the four classes to be modeled in the simulation, a fare index, indicating the relative fare levels within any given market, was developed from the fare ratios in the surveyed markets. From a base of 100 for the Y class fare, lower-yield classes were assigned a relative fare level. These relative price points did not precisely match the collected fare data in all markets. However, the chosen index values are at least representative of the average price / quality trade-offs implied in the domestic air transportation marketplace for the fare products defined here. The Q class fare is used as the base for all calculations involving attributed cost behavioral ratios (Section 3.2).

The final information presented in Table 4.1 covers the costs attributed to each restriction category by the two passenger types. Market research conducted at Boeing (1988) serves as the primary justification for these estimates. In the Boeing stated preference study, members of three passenger segments—nondiscretionary business, discretionary business, and pleasure—were asked about their willingness to travel, in markets exceeding 1300 miles, under various levels of advance purchase and minimum stay restrictions, as well as the difficulty of complying with a Saturday night stay requirement. Responses were expressed as the percentage of each group able to meet each package of restrictions.

A rough translation of these percentages to attributed dollar costs can be performed fairly easily, due to assumptions in the demand generation stage of PODS. Specifically, attributed costs are taken to be normally distributed within a passenger type population. The inputs which govern the restriction cost distribution are the individual mean values and a common cost k-factor which generates variability about the means. Working from the established fares for the products to which the restrictions apply, an attributed cost can then be derived which reflects the behavior indicated by the percentages of the original Boeing study.

As a stylized example, imagine there are only two fare products, priced at $60 and $50, and that the only distinguishing feature is the existence of a Saturday night stay restriction for the cheaper option. Further, the Boeing research reveals that 50% of the surveyed business travelers are willing and able to sacrifice their weekend plans and stay over a Saturday night, given a reasonable incentive. In these circumstances, attribution of a $10 cost (for the average business passenger) to the Saturday night stay requirement would reproduce the preferences implied by the survey. Half of the business passengers would value their free weekend at more than $10 and refuse the trade-off presented by the $50 product, while the other half would find the $10 savings sufficient to compensate for the Saturday night stay. If the percentage were lower (higher), the mean attributed cost would be revised upward (downward).

An important consequence of this restriction cost arrangement is that it allows passengers of either type to travel in any available fare class which is competitive on the basis of *generalized cost*—fare modified by attributed restriction costs. Nondeterministic restriction cost valuations create the possibility that some will be able to "jump the fences" designed to map perceived passenger types to fare classes. This complication, while

Fare Code	Dollar Price	Advance Purchase	Restriction Categories and Individual Costs								Total Average Attributed Cost	
			Saturday Night Stay		Round-trip		0% Refundability		50% Refundability			
			Business	Leisure	Business	Leisure	Business	Leisure	Business	Leisure	Business	Leisure
Y	$100	-	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	$0	$0
B	80	3	N/A	N/A	$15	$0	N/A	N/A	$5	$10	20	10
M	50	7	$25	$5	15	0	$10	$15	N/A	N/A	50	20
Q	40	14	25	5	15	0	10	15	N/A	N/A	50	20

N/A = restriction not applicable to fare product

Table 4.1. OCSE Standard Fares, Restrictions, and Costs

discouraging for the marketing manager seeking perfect market segmentation, is a definite asset from a modeling perspective, since the fallibility of fences is a proven industry reality. Because diversion can work both ways, authenticity also brings a potential reward by permitting leisure passengers to "buy up" to higher-yield products, should the discount options be unavailable. In the base behavioral profile, the acceptable price ratio has been chosen so that the average vacationer will consider all options, except the full-fare Y class product. Section 4.3 examines the sensitivity of this assumption.

To put it plainly, where previous simulation attempts have rested on the existence of independent demand by fare class, PODS takes the more defensible view that fare class choice, for an individual passenger, depends on a subjective appraisal of the prices and restrictions associated with each class. Left to their own devices, passengers may well not book in their "proper" class, making the RM task a more challenging one and likely limiting the revenue gain achievable with RM under a realistic booking process.

With only one measure of variability, or degree of freedom, in the specification for all restriction cost distributions, the survey percentage results cannot be jointly reproduced, in general. Indeed, such a mathematical exercise seems excessive for the present purposes. Rather, the market research findings have been used to guide the selection of suitable attributed costs through repeated directed simulation runs in which costs were varied and the resulting fare class traffic distribution reviewed—pointing the way to the next estimate of attributed cost. After approximate values were found for the Saturday night stay restriction, subjective judgment on the relative severity of the other restriction categories generated the remaining attributed costs for use in OCSE.

Although it was deemed desirable to develop a behavioral framework in which passengers attribute costs to individual restrictions, PODS does not in fact consider these constituent costs in the final analysis. Only the sum of the applicable restriction costs influences the evaluation of path / fare class preference. Consequently, the sensitivity ranges of Section 4.3 pertain to the total cost assessed on each affected fare class, rather than the allocation among the refundability, roundtrip, and Saturday night stay restrictions.

4.2 Further Operational Definitions

Beyond the fare product structure, several other pieces of information are included in OCSE. These can be divided into two broad classes: items which further define the booking process and market setting, and simulation parameters.

The sample booking curves of Figure 2.1, illustrating the trend in cumulative reservation requests as the day of departure approaches, are those adopted for the simulation. They have been distilled from actual aggregated carrier data and checked against similar tables compiled for a recent MIT report on demand driven dispatch (Waldman, 1993, p. 46). The number of booking periods has been set at eight so that there are several periods in which all fare products are available (or more precisely, not closed due to advance purchase requirements).

It will be noticed that the two dimensions along which the curves differ are passenger type and market length of haul. The first of these requires little elaboration—one of the hallmarks of a leisure passenger is the ability to make travel plans (and hence reservations) relatively early in the booking process, shifting the cumulative probability distribution towards the northwest corner of Figure 2.1. However, length of haul also influences the booking trend, because a more distant destination is often associated with a longer stay and, in turn, a longer planning horizon. Conversely, short haul, shuttle-style flights may be taken with less deliberation. Within a passenger type, the cumulative booking distribution for those traveling in a short haul market will thus be moderately weighted towards the day of departure, compared to long haul flights. Figure 2.1 indicates that the distinction between passenger types outweighs the length of haul factor (e.g. the short haul business curve more closely resembles the long haul business curve than the short haul leisure curve). The long haul threshold has been set at 1000 miles. Long haul markets have also been chosen for all competitive experiments. The sensitivity analysis considers a short haul market and finds no appreciable difference in the ultimate traffic statistics.

All stochastic processes assumed to follow a normal distribution—demand generation and restriction cost attribution—have been assigned a k-factor of 0.3.[34] This

[34]This has been set directly in all one market scenarios. However, recalling the nested k-factor arrangement described in Section 3.2, the total variability in multiple market scenarios has been evenly

value has repeatedly surfaced in empirical studies of air transportation demand and no other ratio seems a better candidate for the default setting. The sensitivity analysis of the next section explores the consequences of departing from this standard. The z-factors, which introduce correlation among booking periods, have been set at 4 and 2, creating moderate correlation. These selections are not tested below but alternative values would be expected to have only a minor effect on the overall booking results. An additional demand generation parameter describes traffic composition by passenger type. The base ratio for leisure to business volume is 2:1, reflecting a mixed market. Other potential splits are evaluated in the sensitivity analysis.

Seed values for the mean leg demand forecasts, to be consulted until a significant booking history has accumulated, have been obtained by applying a formula to the specified market demands, in all markets traversing the leg, for the two passenger types. Estimates of forecasting error follow directly from the means through the use of a constant k-factor, again equal to 0.3. In informal tests, the weights used in the mean forecast formula (which would also alter the error) were found to have no noticeable effect on the average booking statistics. More critical to the early forecasting stages is the accurate portrayal of *attributed fares*, which are the fare class revenue values used in the EMSR calculations of Section 2.2.2.1. PODS offers two methods of computing these attributed fares (Section 3.4), which can be toggled independently for each participating carrier. In all scenarios, OCSE adopts the simpler nonstop attribution method. After implementation of partial network-control, the weighted average may be more appropriate.

Other airline-specific method indicator inputs affect the RM forecasting submodel. Estimation of the booking parameters—the exercise of building, from historical reservations data, an internal representation of the underlying unconstrained booking curves—is performed using exponential smoothing of the untruncated observations. As a user-defined option, truncated observations, or those which closed prior to the day of departure, may also be incorporated in the forecasting process after detruncation (Section 3.4). Because exclusion of the truncated data would seriously bias demand forecasts, all observations—both untruncated and detruncated—are used. The basic pick-up forecasting model of Chapter 2 has been universally adopted.

---

spread among system, market, and passenger type component k-factors, with each taking a value of 0.17 (refer to Equation [3.1])

No experimentation with carrier image factors has been attempted. While PODS permits modeling of carrier preference on a market-by-market basis (to capture strength in hub operations, for instance), such detail is beyond the scope of this thesis. Similarly, all aircraft capacities have been standardized at 100 seats. Finally, OCSE accepts the default PODS settings for mean schedule tolerance (related to decision window construction) and time-of-day demand distribution, developed from Boeing market research. These areas certainly merit future study, as the research focus turns from the modeling of stylized, exemplar scenarios to simulation of the competitive landscape in actual markets.

Section 3.1 defines the simulation term, *observation*, as one departure day. However, the true sampling unit is a set, or time series, of observations, which will be referred to as a *trial*. There are two reasons why an observation is not a particularly meaningful unit of analysis. First, operating performance for a single departure is subject to variation introduced in the demand generation stage. Isolated observations, then, may be outliers unrepresentative of the average result and inappropriate for comparative inference. Second, the effect of RM forecasting and capacity allocation can only be seen after a relevant booking history has been compiled. A lone observation must rely on the initial mean forecast demands and errors, which, if inaccurate, may lead to a result inferior to that for an unmanaged, first come, first served process.

By choosing a trial as the sampling unit, rather than an isolated observation, both of these objections are addressed. First, revenue and traffic summary statistics, for repeated observations ruled by the same demand parameters, will be more stable than samples of individual observations. Furthermore, the set of observations constitutes a continuing history used to populate the RM booking parameter estimation procedure. With the conceptual advantage of a trial established, three sizing inputs must be specified.

The first concerns the number of observations to be stored in the forecaster. Deciding to keep more observations will lower the chance that a single forecast will be made on the basis of unusual recent booking data, but also increase the time it takes for an outlying observation to "wash out" of the RM system. Of the two elements in this trade-off, the second is probably less significant since the exponential smoothing technique places little weight on non-recent terms. Consequently, a value of 20 has been selected, towards the upper end of the range found in industry practice.

The second input requirement arises with the potential errors in the seed forecast inputs, as noted above. In the calculation of summary statistics, it would seem appropriate to ignore, or "burn," early observations which may have been unduly influenced by faulty initial forecasts. Finally, the total observation size of the trial must be defined. The trial should be sufficiently large that a marginal observation would have a small effect on the cumulative averages, but cannot be made arbitrarily large without incurring serious penalties in execution time (for negligible gains in accuracy). After some experimentation with the observation size inputs, and monitoring of the resulting statistics over the course of the trial, the burn and total counts have been chosen as 200 and 2000, respectively.

Though most of the sample size issues have now been considered, one unanticipated result remains to be discussed. During preliminary model validation with scenarios consisting of two identical paths (the class of scenarios studied in Section 5.1, with two carriers practicing inventory control, each with one departure), the observed evolution of the revenue statistics seemed puzzling. Early in the trial, one carrier would often take a sizable revenue lead. In itself, this is entirely consistent with chance fluctuations in the passenger assignment process—faced with an identical product available on either carrier, an indifferent passenger chooses one at random. However, this revenue advantage did not dissipate as quickly as expected. After some reflection, it became clear that the cause for the persistent disparity lay in the correlation of the observations, created by the intervention of the forecasting and inventory control systems.

The anomaly begins before a meaningful booking history has been recorded, with a random traffic imbalance between the carriers, as described above. The recipient of unusually high bookings proceeds to adjust his future booking limits and protection levels in anticipation of continuing patronage, setting aside more seats for late arriving business passengers. Conversely, the unfortunate rival expands discount availability, seeing little potential for RM in such a seemingly sparse market. Neither carrier knows that the recent performance is not indicative of the true, balanced state of affairs. Nevertheless, a sort of "self-fulfilling prophecy" ensues for the next several departures. The carrier who felt the demand spike tends to maintain his revenue advantage, even under normal demand conditions, because business passengers can be accommodated (there are just fewer than expected). However, the rival expecting the worst necessarily finds it—without reserving space for late bookings, leisure passengers are allowed to "crowd out" potential late-arriving business traffic. Eventually, the imbalance corrects itself, either by an offsetting random swing in the opposite direction, or once enough typical observations have entered the forecaster database. Yet the chance demand disturbance may still appear in the final aggregate statistics.

This potential sensitivity to initial demand conditions has prompted the use of multiple independent trials for each competitive scenario. In the experiments of Chapter 5, five to ten trials have been run and collapsed to obtain "averages of the average trial statistics," which should be nearly free of the random effects carried through the correlated observations of a single trial.

4.3 Sensitivity Analysis of Noncompetitive Simulation Attributes

To assess the sensitivity of OCSE to the assumptions described in the first part of this chapter, two or more alternative values for each attribute have been substituted (in turn) for the base value, and the resulting statistics are presented here. All attributes have been tested in isolation, though the direction and approximate magnitude of interaction effects can be determined from the component attribute changes. These tests are performed in a one path environment to avoid the confounding effects of competition introduced in the experiments of Chapter 5, all of which apply the base behavioral and demand composition assumptions identified here.

Two levels of *unconstrained demand factor* (DF), defined as the ratio of unconstrained demand to aircraft seats, are included in the analysis. The lower value, 0.9, lies at the low end of the DF range over which RM would be expected to have a noticeable revenue impact. Capacity is considerably more strained at the second DF of 1.2—for a typical departure, the number of booking requests exceeds the available seats by 20%. Because the higher DF level is likely to show greater variation under the different sets of assumptions, all of the traffic and revenue comparison graphs below are based on a DF of 1.2. Having established these upper bounds, written remarks and a set of summary tables (4.1-4.4) note the observed sensitivities for the less extreme demand factor.

Because the sensitivity for certain attributes will clearly be influenced by the presence of inventory control, this factor has been treated separately in all attributes. First come, first served (FCFS) and EMSRa graphs are shown on facing pages.

The tested attributes can be divided into two classes:

- passenger behavior, including:
  - acceptable price ratio (for leisure passengers)
  - advance purchase restrictions
  - average restriction cost (for business passengers)
  - attributed cost k-factor

- demand composition, including:
  - system-level demand k-factor
  - passenger type split

4.3.1 Passenger Behavior

Acceptable Price Ratio (APR) (Leisure Passengers) [Figures 4.1-4.2]

Expected direction of effect: A larger ratio would cause revenue to rise, as leisure passengers are more willing to buy up to higher-yield fare products. As noted in Section 3.3, available fare products are dismissed out of hand if they fail to meet the acceptable price criterion. Consequently, APR may be among the more influential of the behavioral assumptions.

Base attribute setting: ratio set so acceptable price for average leisure passenger matches class B fare.

Alternative settings for average acceptable price:

[a] equal to class M fare;

[b] greater than class Y fare.

Sensitivity findings: Figure 4.1 clearly demonstrates the shift in the fare class distribution from Q to the adjacent B and M classes, as the acceptable price ratio increases ([a] to [base] to [b]), under EMSR control. The percentage revenue changes from [base] to [a] and [b] are -16% and +18% respectively, at a DF of 1.2, and roughly one-third of these amounts at the more moderate DF of 0.9. FCFS does not derive a significant benefit from the increased willingness to pay, at either DF.

Advance Purchase Restrictions [Figures 4.3-4.4]

Expected direction of effect: Removal of advance purchase restrictions would increase bookings in the affected discount fare classes, thus lowering market revenue, to the extent that these passengers would have otherwise sought accommodation in higher classes. The strength of this shift depends on the fraction of late-arrivers in the market, determined from the cumulative booking curves. Because the total fraction in the last three booking periods is small for the curves employed here, even in short haul markets, the elimination of advance purchase requirements may have only a minimal effect on total revenues.

Base attribute setting: classes M and Q have advance purchase requirements of seven and fourteen days, respectively.

Alternative settings:

[a] only class Q requires a fourteen day advance purchase;

[b] no class requires advance purchase.

Sensitivity findings: Comparing Figures 4.3 and 4.4, under alternative [b] at a DF of 1.2, the results for EMSR and FCFS are virtually identical. With nothing now distinguishing classes M and Q (originally, both have the same restrictions except for advance purchase levels), except a ten dollar difference in price, it is not surprising that the former receives almost no bookings. In alternative [a], EMSR control redirects late-arriving passengers from Q to Y, B, and M classes, with the last seeing the largest absolute change as the only discount product open in the last two weeks of the booking period. The load between M and Q nearly balances for [base] and, as Table 4.2 reveals, between two and three of the rare late-arriving leisure passengers spill on an average departure, because they cannot afford the remaining open options. The EMSR percentage revenue changes at a DF of 1.2 are twice as great as those for a DF of 0.9. The magnitude of the revenue effects belie the *a priori* argument above—apparently, with the Y class fare two and a half times greater than the Q class fare, advance purchase restrictions have substantial revenue consequences, even in markets where most passengers book early.

DF has little effect on the FCFS sensitivities, but a drop in revenue of about 10% for alternative [b] over [base] is significant. Essentially, advance purchase restrictions function as a rudimentary form of control for the unmanaged carrier.

Average Restriction Cost (ARC) (Business Passengers) [Figures 4.5-4.6]

Expected direction of effect: As business passengers assign higher costs to the restrictions attached to discount fare products, the potential for dilution falls. Revenues improve as a larger fraction are captive to the high-yield Y and B classes.

Base attribute setting: average ARC total is equal to class M fare. This creates a situation in which the average business passenger's generalized cost for M class is equal to the unrestricted Y class fare.

Alternative settings for average ARC total:

[a] less than class Q fare;

[b] equal to class B fare.

Sensitivity findings: Figure 4.5 shows ARC to be one of the less influential behavioral factors in the simulation. In the progression from [a] to [base] to [b], there is the anticipated shift in traffic distribution from M, the more expensive discount product, to Y and B classes. Yet the volume of this shift does not generate very different revenues—within ±10% under both DFs. It appears that the management of business passenger allocation already provided by EMSR control means little marginal revenue improvement occurs when ARC rises.

Indeed, the relative effect of varying ARC is more pronounced for the FCFS case (Figure 4.6), particularly from [base] to [b]. In part, this can be traced to the larger incremental revenue earned by redirecting business passengers out of Q class, rather than M class, out of which the movement in EMSR occurs. More generally though, ARC percentage revenue sensitivity runs closer to 15% under both DFs simply because the base FCFS fare class distribution, or *mix*, is so depressed.

Attributed Cost K-Factor (ACKF) [Figures 4.7-4.8]

Expected direction of effect: uncertain. When attempting to understand how increased variability in a process will affect the outcome, it is often helpful to consider asymmetries in the effective range of values which would result in the more stochastic environment.[35] However, this is quite difficult to do here, for two reasons. First, two opposing passenger assignment forces are influenced by ACKF, via the previously examined attributes of maximum acceptable price and attributed restriction cost. The first governs how likely leisure passengers are to buy up, the second how likely business passengers are to find discount restrictions tolerable. The net effect of these two factors under different levels of variability is not immediately apparent.

Furthermore, even isolating one of these factors, say acceptable price, the bounds on the range of attribute values are fairly balanced. From the mean, set equal to the B class fare, the maximum change which has a distinct effect on passenger behavior is +$20, bringing acceptable price to the highest fare, Y class—at this point, and everything beyond, leisure passengers are willing to contemplate all available fare products. On the "down side," once acceptable price drops by $40 (or the ratio equals zero), leisure passengers must book in Q class. By this crude analysis, a higher k-factor, which creates an opportunity for more frequent and extreme departures from the mean, might bring higher revenues because the best "up side" result at +20 would appear more often than the worst case -40. This argument ignores the relative payoffs at each end of the acceptable price spectrum, as well as the confounding factor of inventory control, but does introduce the complexities in predicting even the direction of the expected effect.

Base attribute setting: 0.3.

Alternative settings:

[a] 0;

[b] 0.5.

Sensitivity findings: Figure 4.7 and Table 4.4 indicate that the revenue change produced when moving from [base] to the deterministic case [a] is a significantly positive one under EMSR (+15%/+10% with DF of 1.2/0.9). M class traffic drops precipitously, shifting to the upper classes, and there is a net load gain of over two passengers at a DF of 1.2. Increasing ACKF to [b] has a comparatively minor negative effect.

[35] For a successful application of this approach, see the discussion for System-Level Demand K-Factor below.

In contrast, the FCFS graph (Figure 4.8) at a DF of 1.2, and the results at a DF of 0.9 (Table 4.5) reflect the mixed signals from above. Minimal traffic redistribution has virtually no effect on the average revenue figures. At least the near Y-Q class polarity of the deterministic case [a] may be readily explained—all leisure and early-arriving business passengers (who always value the discount fare savings as greater than the accompanying restriction costs) book entirely in Q class, and the remaining capacity goes to the last-minute requests.

Given that the forces dependent on ACKF are only poorly understood at this point, there is some consolation in the discovery that the sensitivity for higher variability, at least, appears to be minimal (less than 5% in all cases). In estimating the range of intra-type passenger behavior, greater variability seems more probable than less (let alone determinacy).

4.3.2 Demand Composition

System-Level Demand K-Factor (SKF) [Figures 4.9-4.10]

Expected direction of effect: Revenue increases as SKF falls. Although another k-factor test, the outcome here is easier to predict, because of the definite "ceiling" created by aircraft capacity. Extreme values greater than the mean will have little effect on total revenue because a full flight for which twenty passengers are turned away contributes the same revenue as one which turns away fifty passengers. Conversely, the down side potential is nearly unlimited—revenue falls with one empty seat and continues to suffer until the number of reservations reaches zero. A higher SKF raises the chance of low *load factors* (passengers *carried* relative to capacity), without a compensating benefit for large outliers, guaranteeing lower average revenues.

Base attribute setting: 0.3.

Alternative settings:

[a] 0;

[b] 0.5.

Sensitivity findings: A carrier with EMSR control can take good advantage of the deterministic situation in [a], realizing the stability of its business traffic component and revising protection levels accordingly. As would be expected, DF has a significant effect on SKF sensitivity. At a DF of 1.2, the revenue changes for [a] and [b] are +19% and -14% respectively. Under a DF of 0.9 (Table 4.4), these fall to +5% and -7%. An FCFS carrier derives a smaller, but still consequential, revenue gain from increased certainty, due entirely to the higher loads which result. Total load changes for the SKF range are the largest of any attribute test.

Passenger Type Split [Figures 4.11-4.12]

Expected direction of effect: Revenue rises as the business : leisure ratio grows.

Base attribute setting for business to leisure ratio: 1:2.

Alternative settings for business to leisure ratio:

[a] all leisure;
[b] 2:1;
[c] all business.

Sensitivity findings: Tables 4.1 to 4.4 show passenger type split to be the single most important revenue determinant in this sensitivity review, for all control and DF combinations. Starting from [base] at a DF of 1.2, EMSR and FCFS carrier revenues fall by 35% and 12%, respectively, in an all leisure market, in which no inventory management method can avert a meager, charter-flight style yield result. At the other end of the scale, revenues climb by 42-56% in an all business market. FCFS accomplishes the feat entirely through passive traffic redistribution, while EMSR posts the added bonus of eliminating leisure passenger spill with the absence of leisure passengers, thus carrying two or three extra passengers on average. At the lower DF, the percentage effects are just as strong in FCFS and only slightly muted for [a] and [b] under EMSR.

* * *

This chapter starts with a review of the fare product structure, comprising the prices and restrictions attached to each fare product, which has been adopted in OCSE. The values chosen for several other inputs, which affect the booking process and RM environment, are then described. Section 4.3 conducts sensitivity analyses of the attributes governing passenger behavior and demand composition. Supporting graphs and tables demonstrate the revenue and load distribution effects which result under alternate settings.

In order to leave manageable the competitive scenario simulations which follow in Chapter 5, all passenger behavior and demand composition attributes are held constant at their base settings. The worth of RM under various competitive circumstances can then be explored under a single set of representative, intermediate attribute values. The relative strengths of the sensitivity findings demonstrated above, with acceptable price ratio and passenger type split among the most influential attributes, should be borne in mind when reviewing the base case competitive scenario results. An entry in the sensitivity tables may be used as a sort of point elasticity—a rough guideline to extrapolate from the standard revenue and traffic results, after the slight modification of a single behavioral or demand condition. However, direct application of the sensitivity factors to analyses of real-world competitive settings is not recommended. When the modeling of a specific market environment requires large and / or multiple departures from the base settings (creating attribute interactions), the bounds for the applicability of these sensitivity factors have been exceeded. Instead, independent simulation trials should be performed with the new attribute values.

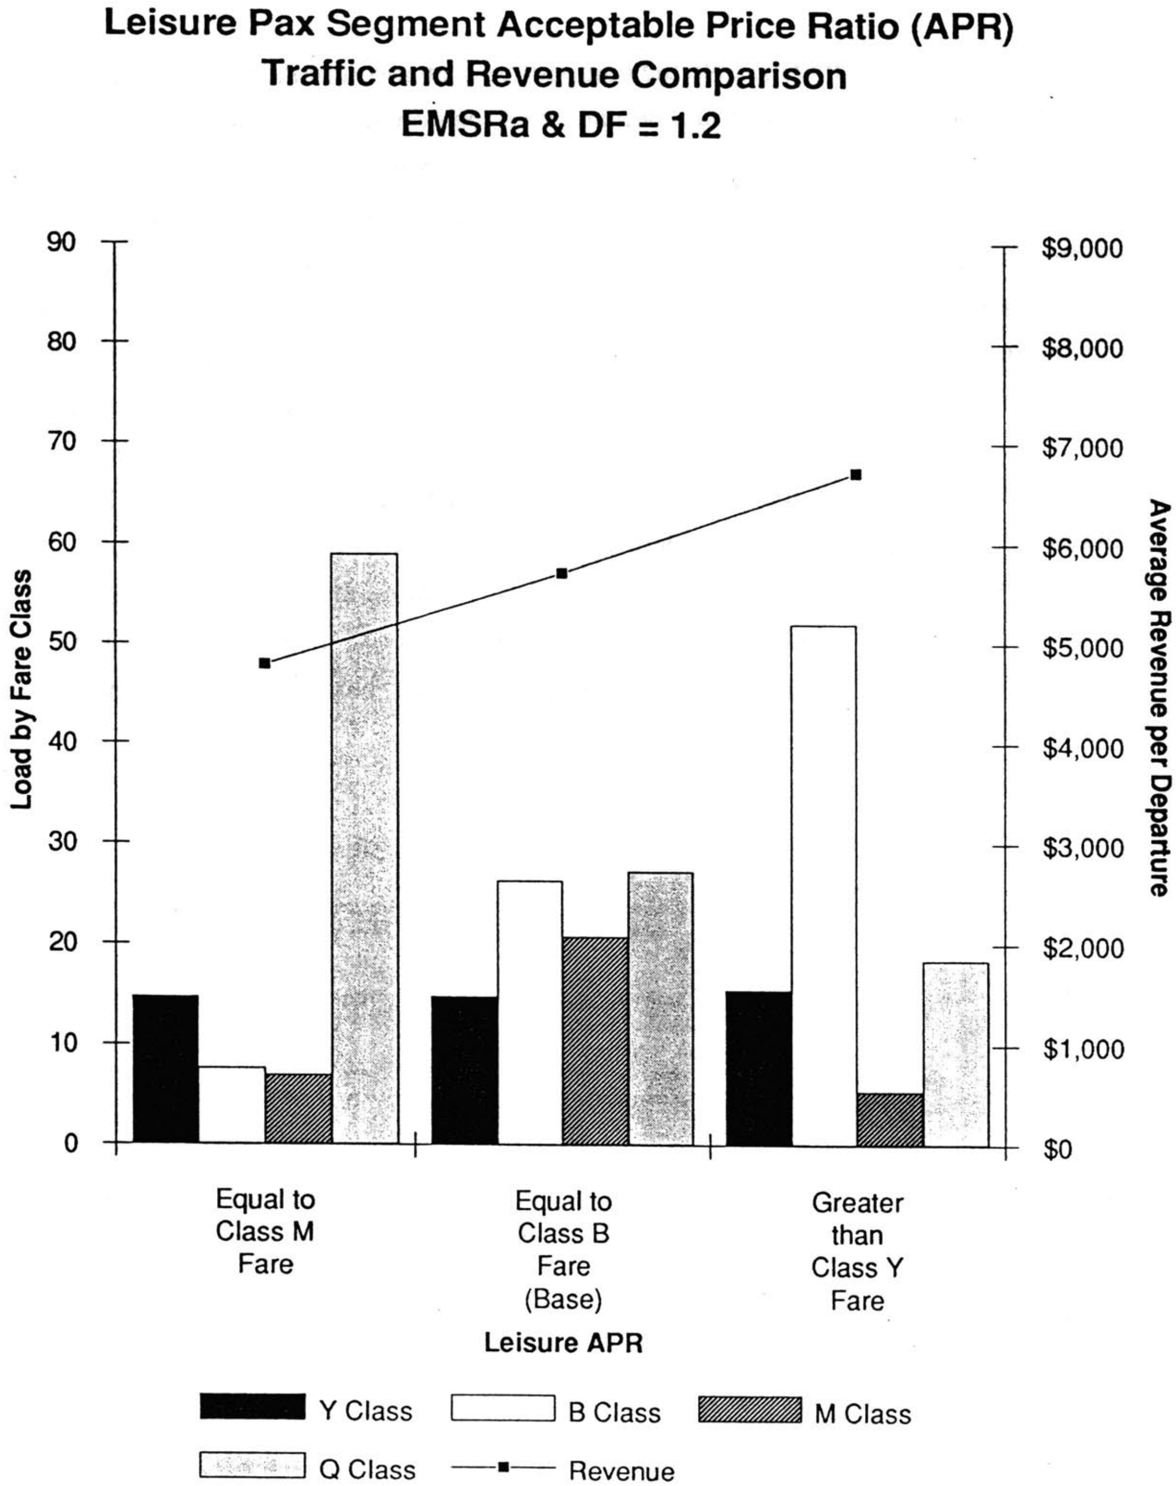

Figure 4.1. Sensitivity of Acceptable Price Ratio Attribute, EMSR

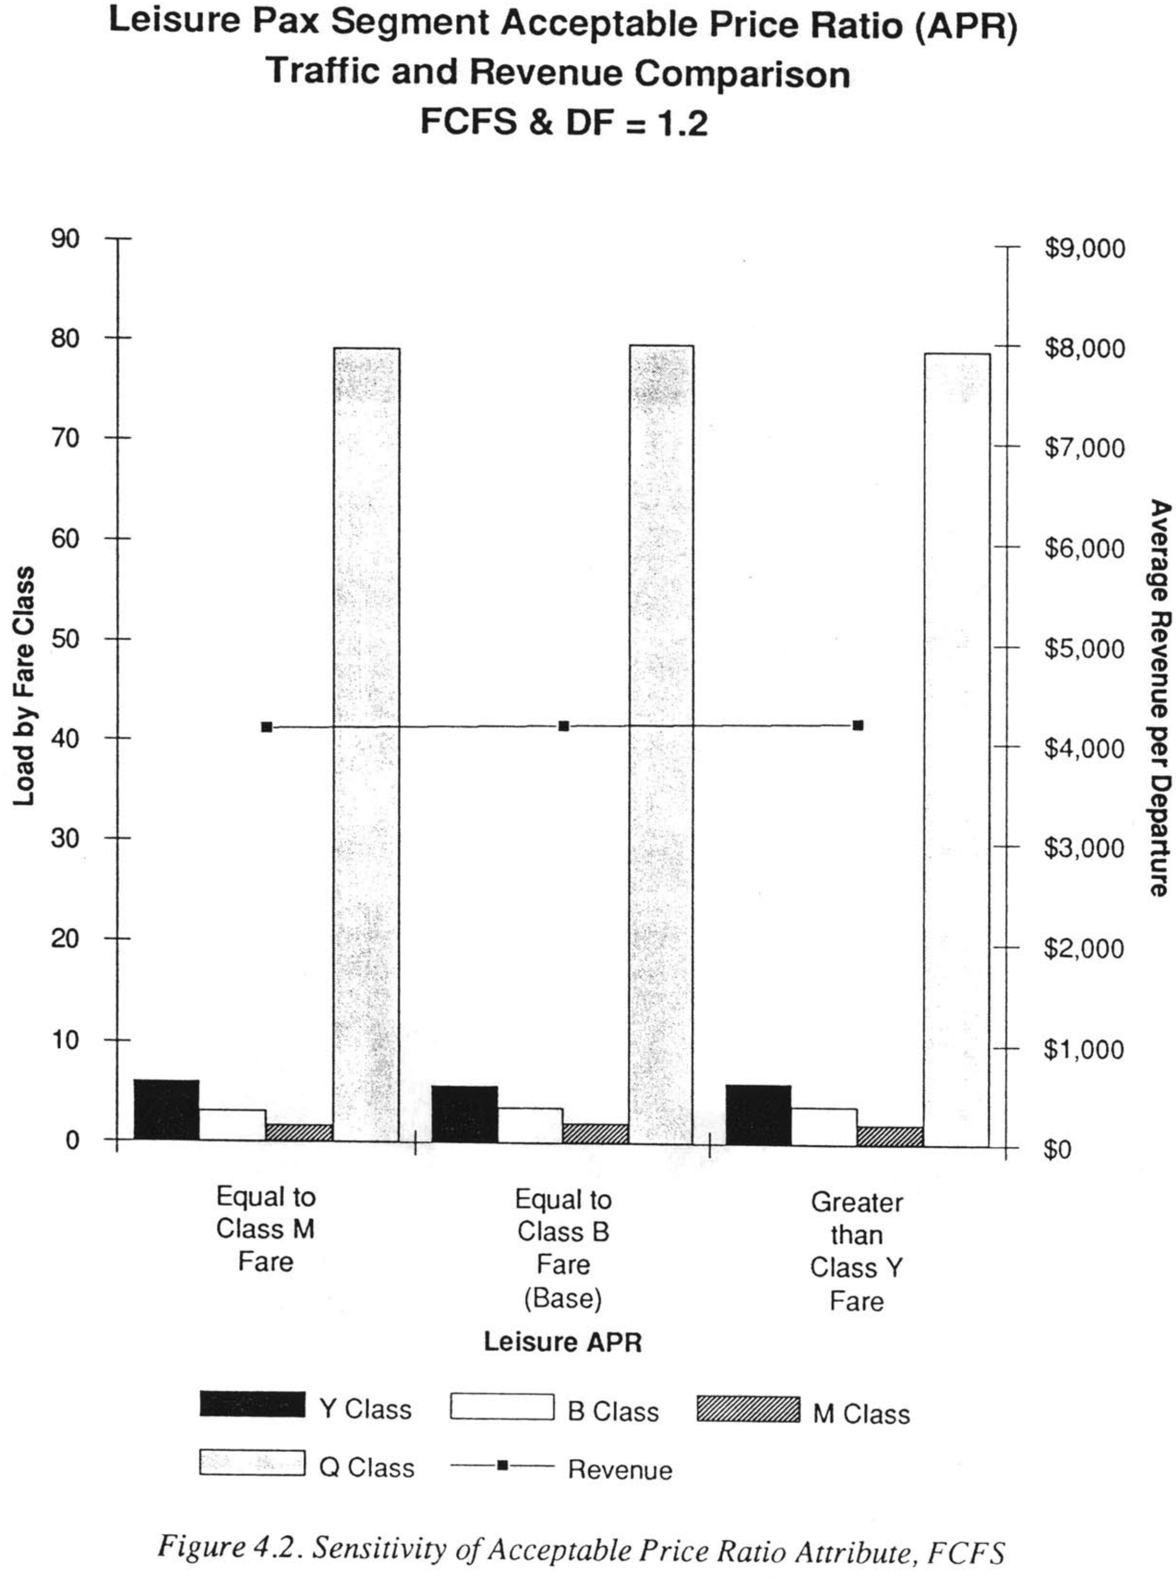

Figure 4.2. Sensitivity of Acceptable Price Ratio Attribute, FCFS

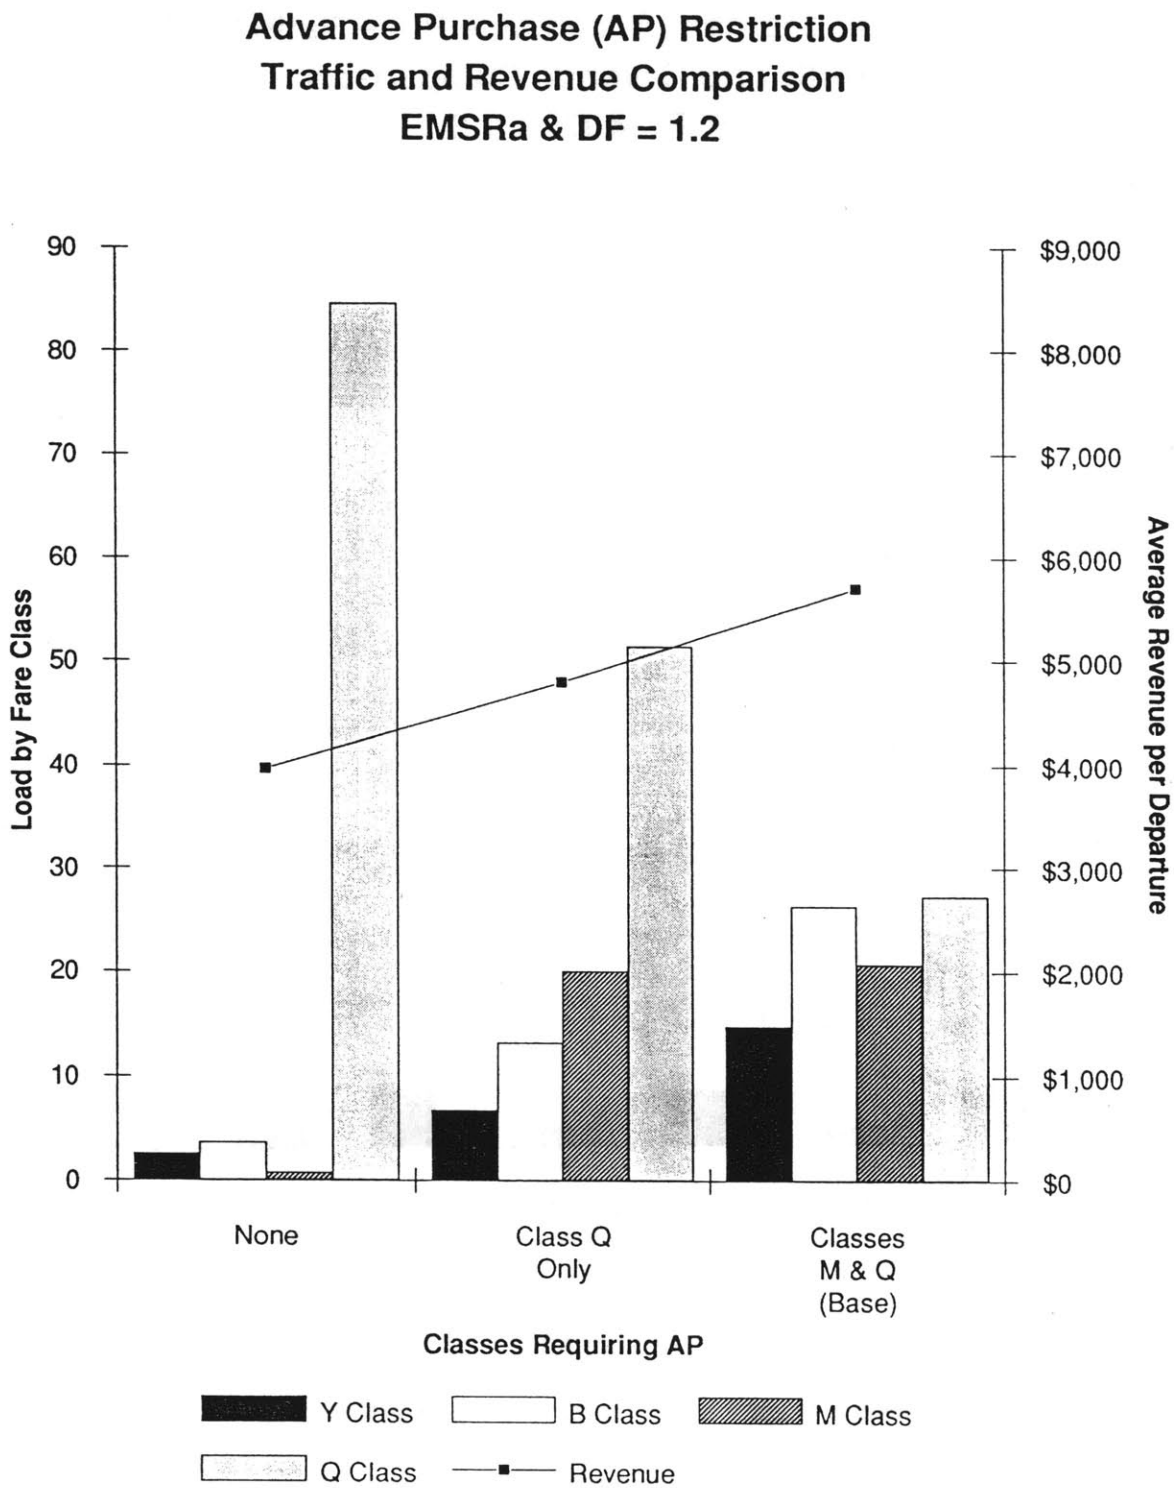

Figure 4.3. Sensitivity of Advance Purchase Restrictions Settings, EMSR

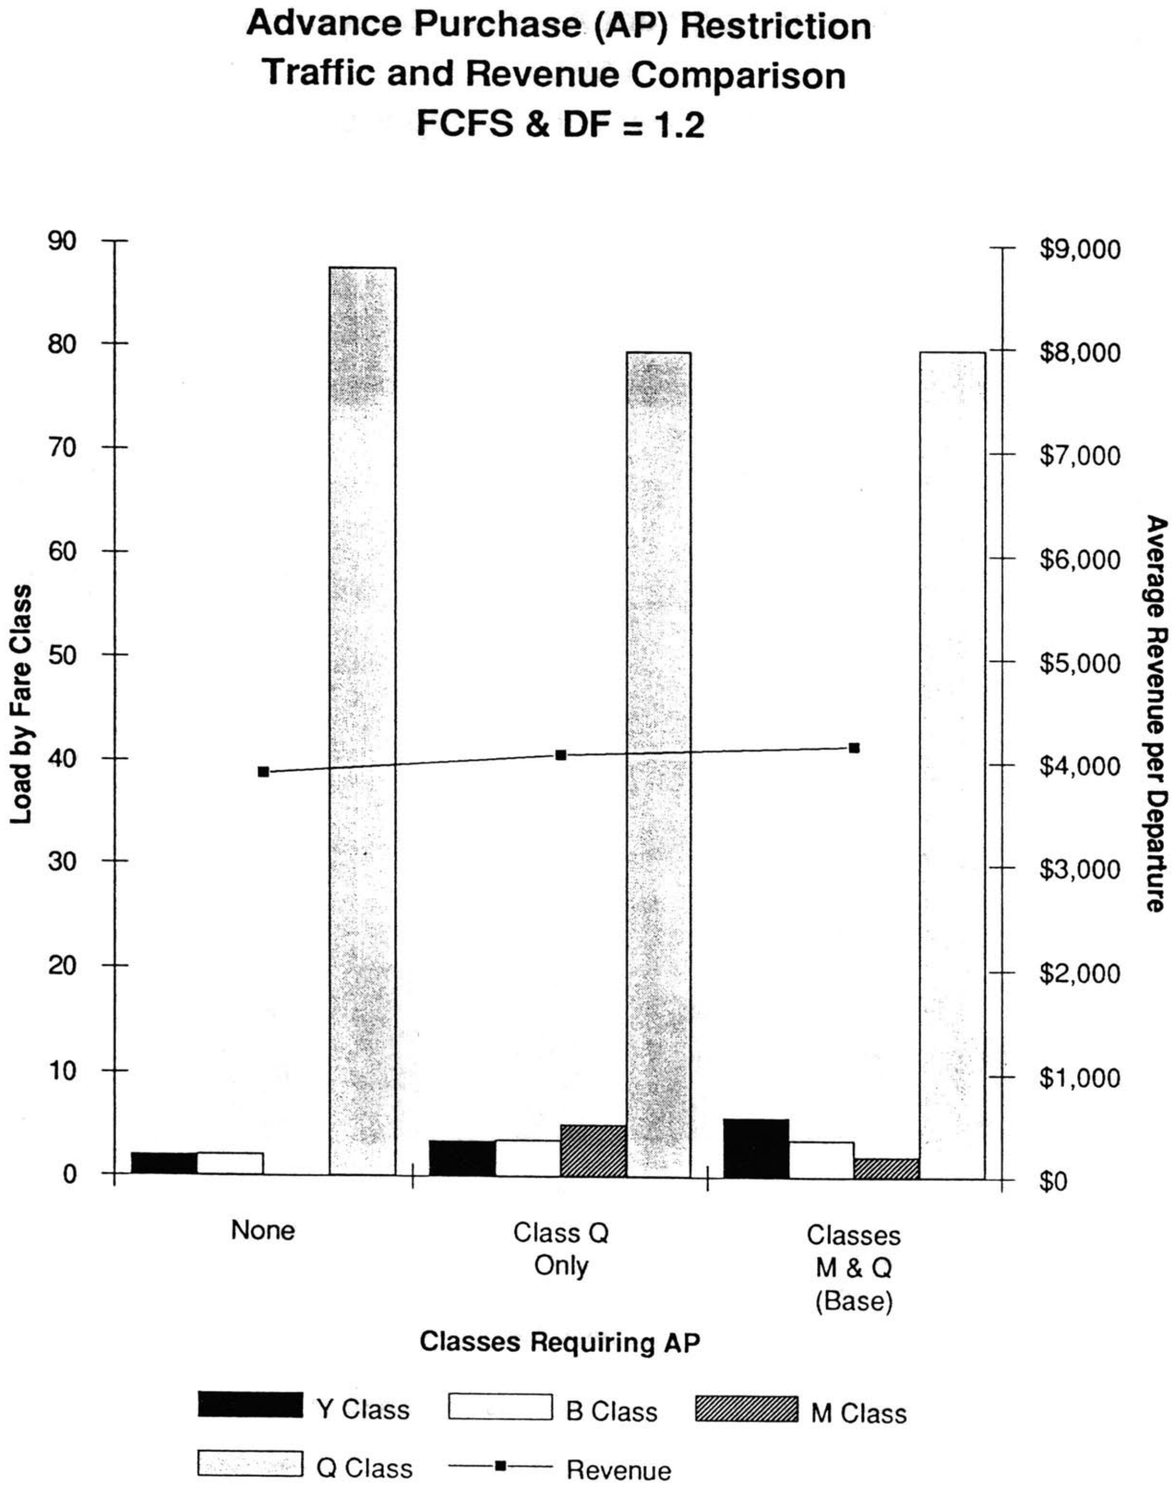

Figure 4.4. Sensitivity of Advance Purchase Restrictions Settings, FCFS

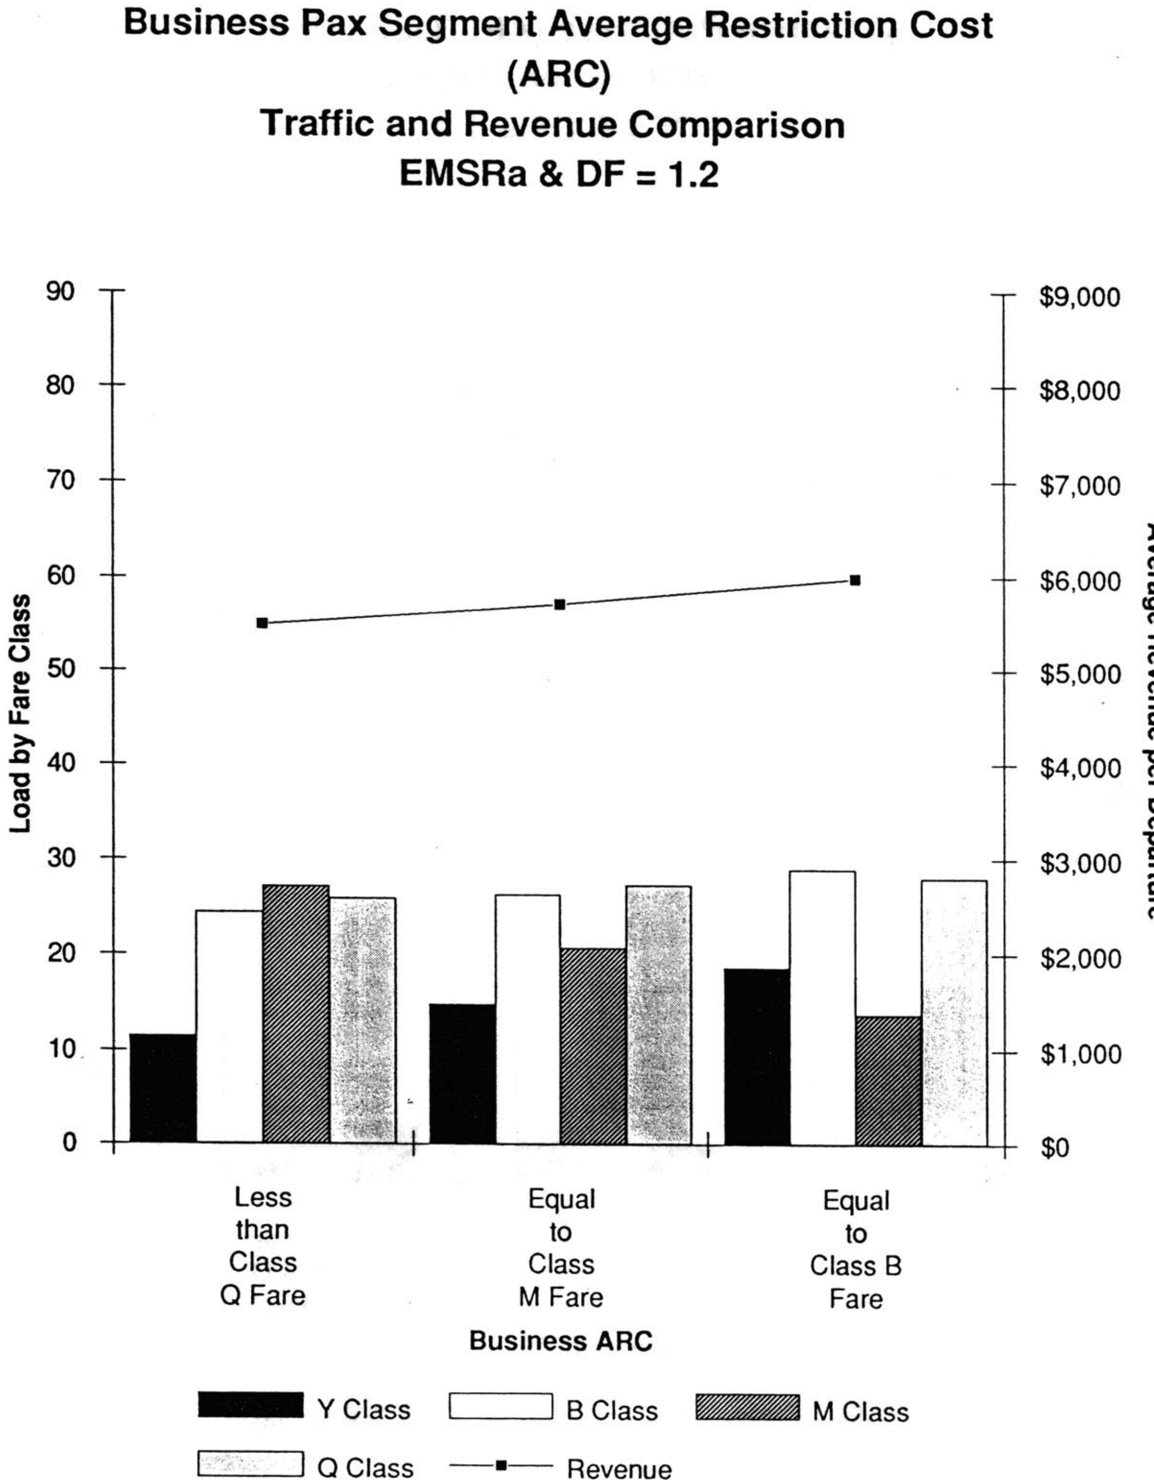

Figure 4.5. Sensitivity of Average Restriction Cost Attribute, EMSR

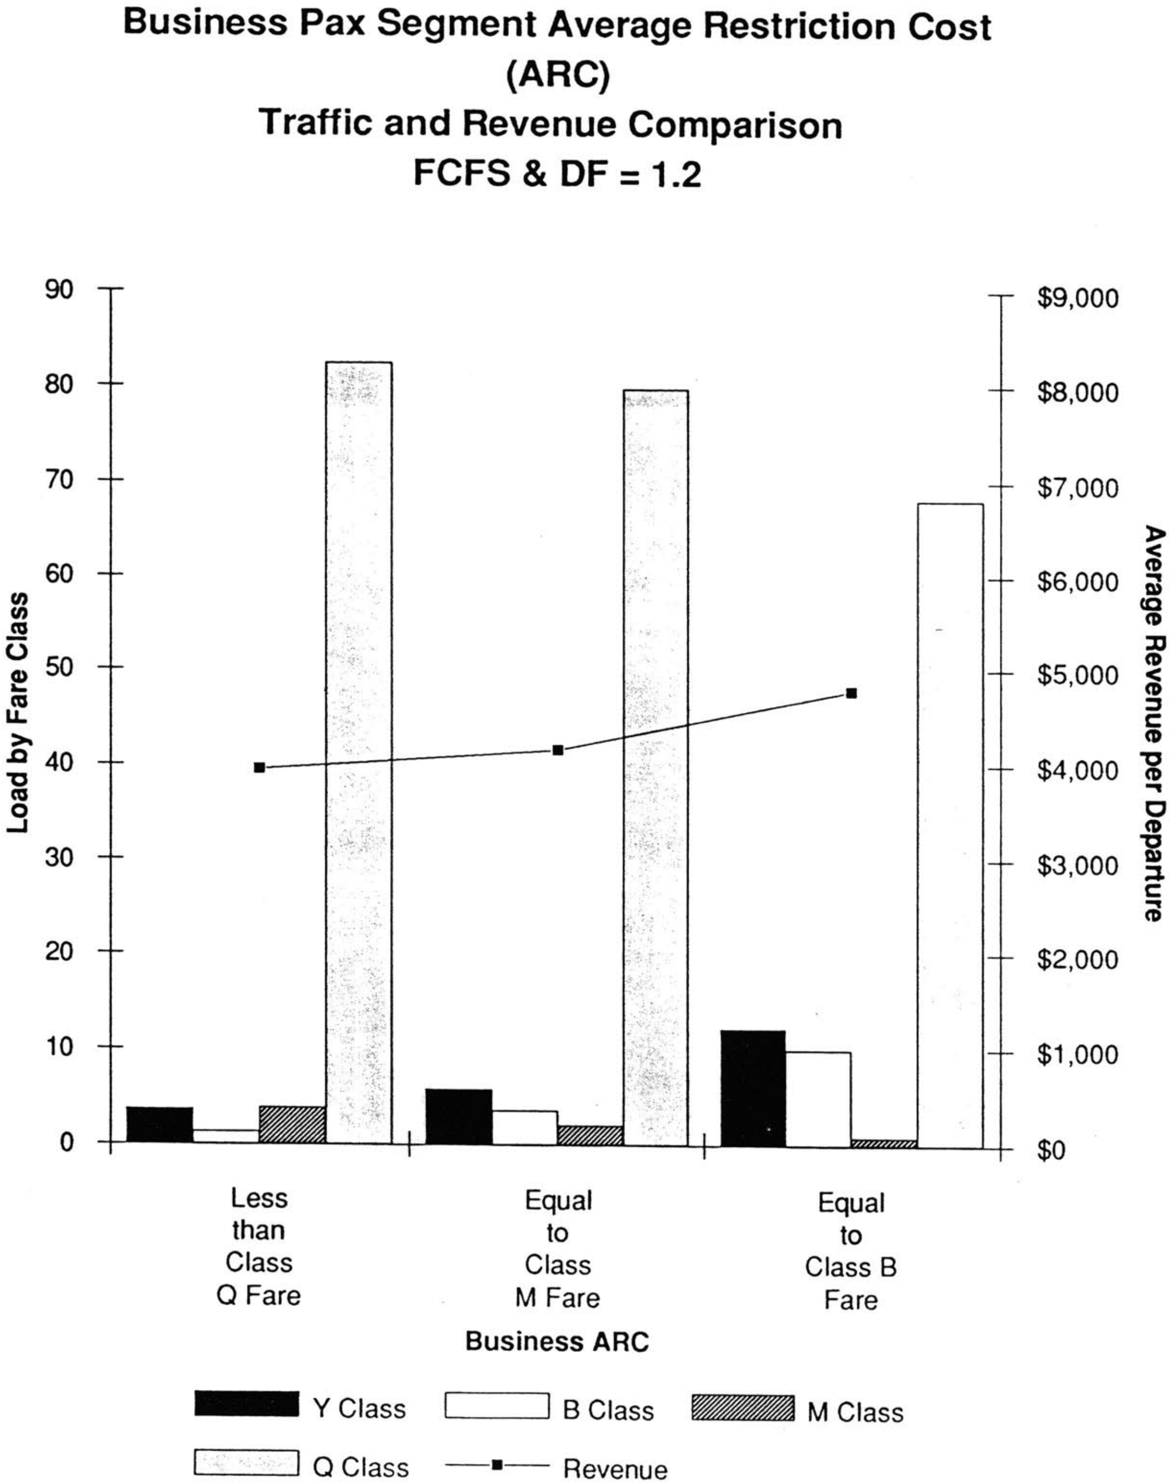

Figure 4.6. Sensitivity of Average Restriction Cost Attribute, FCFS

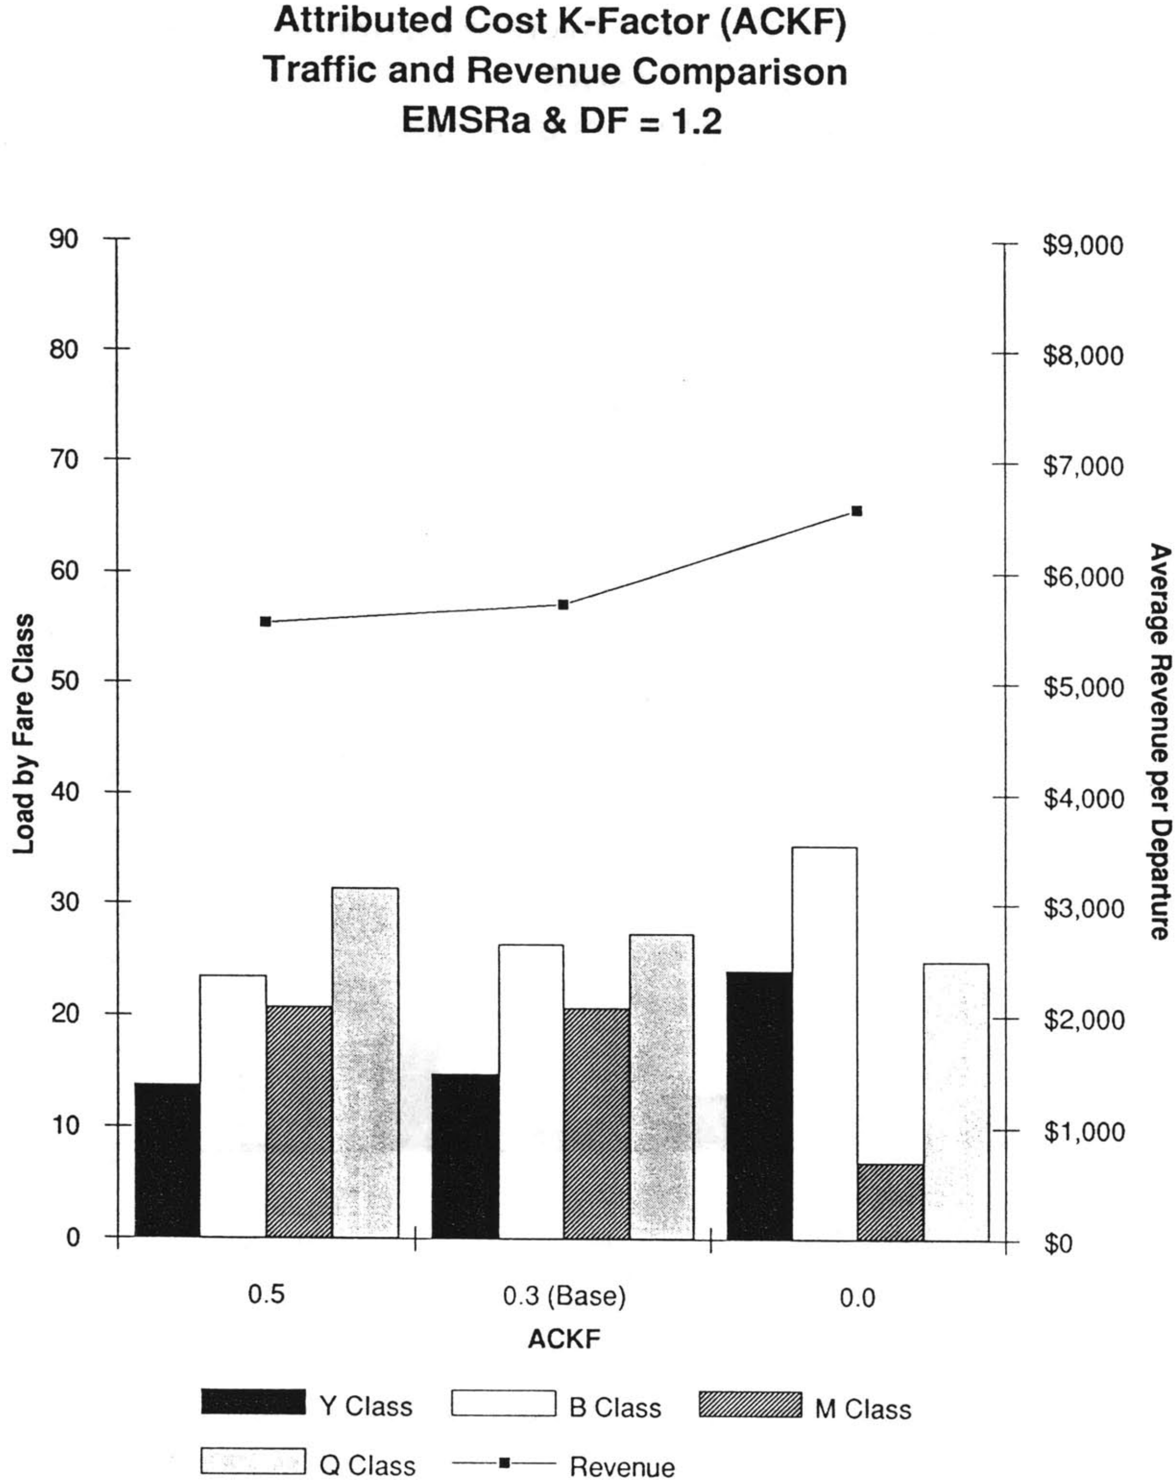

Figure 4.7. Sensitivity of Attributed Cost K-Factor, EMSR

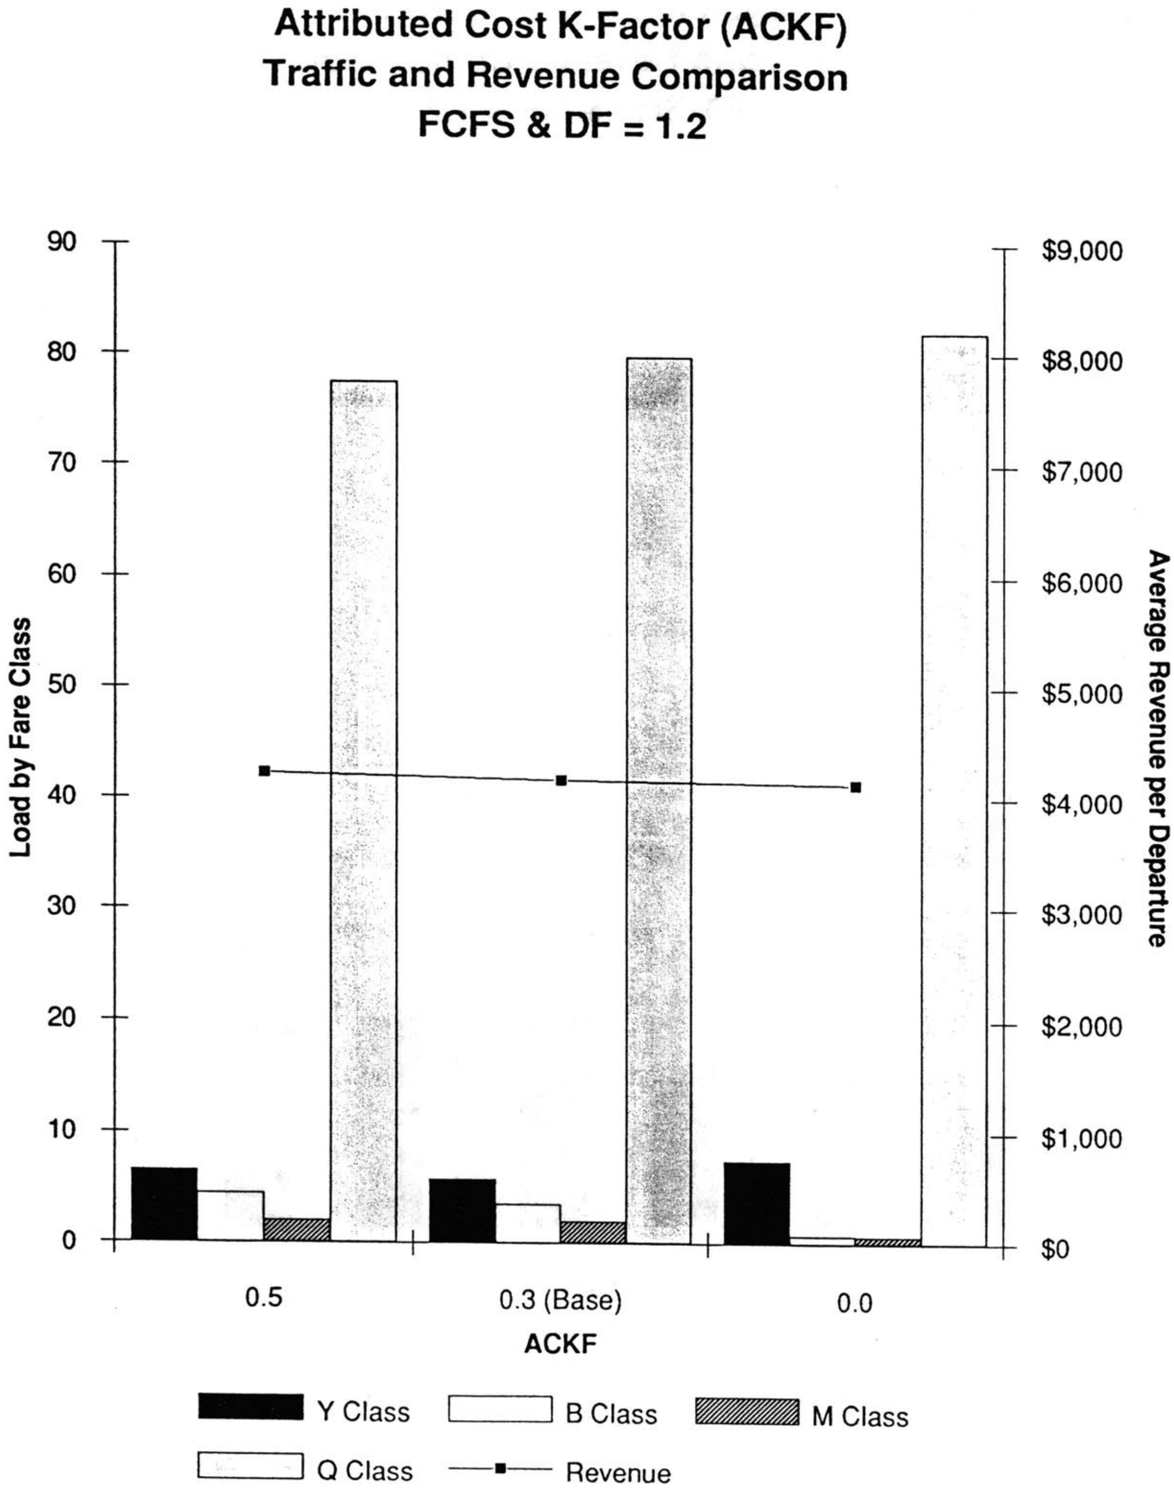

Figure 4.8. Sensitivity of Attributed Cost K-Factor, FCFS

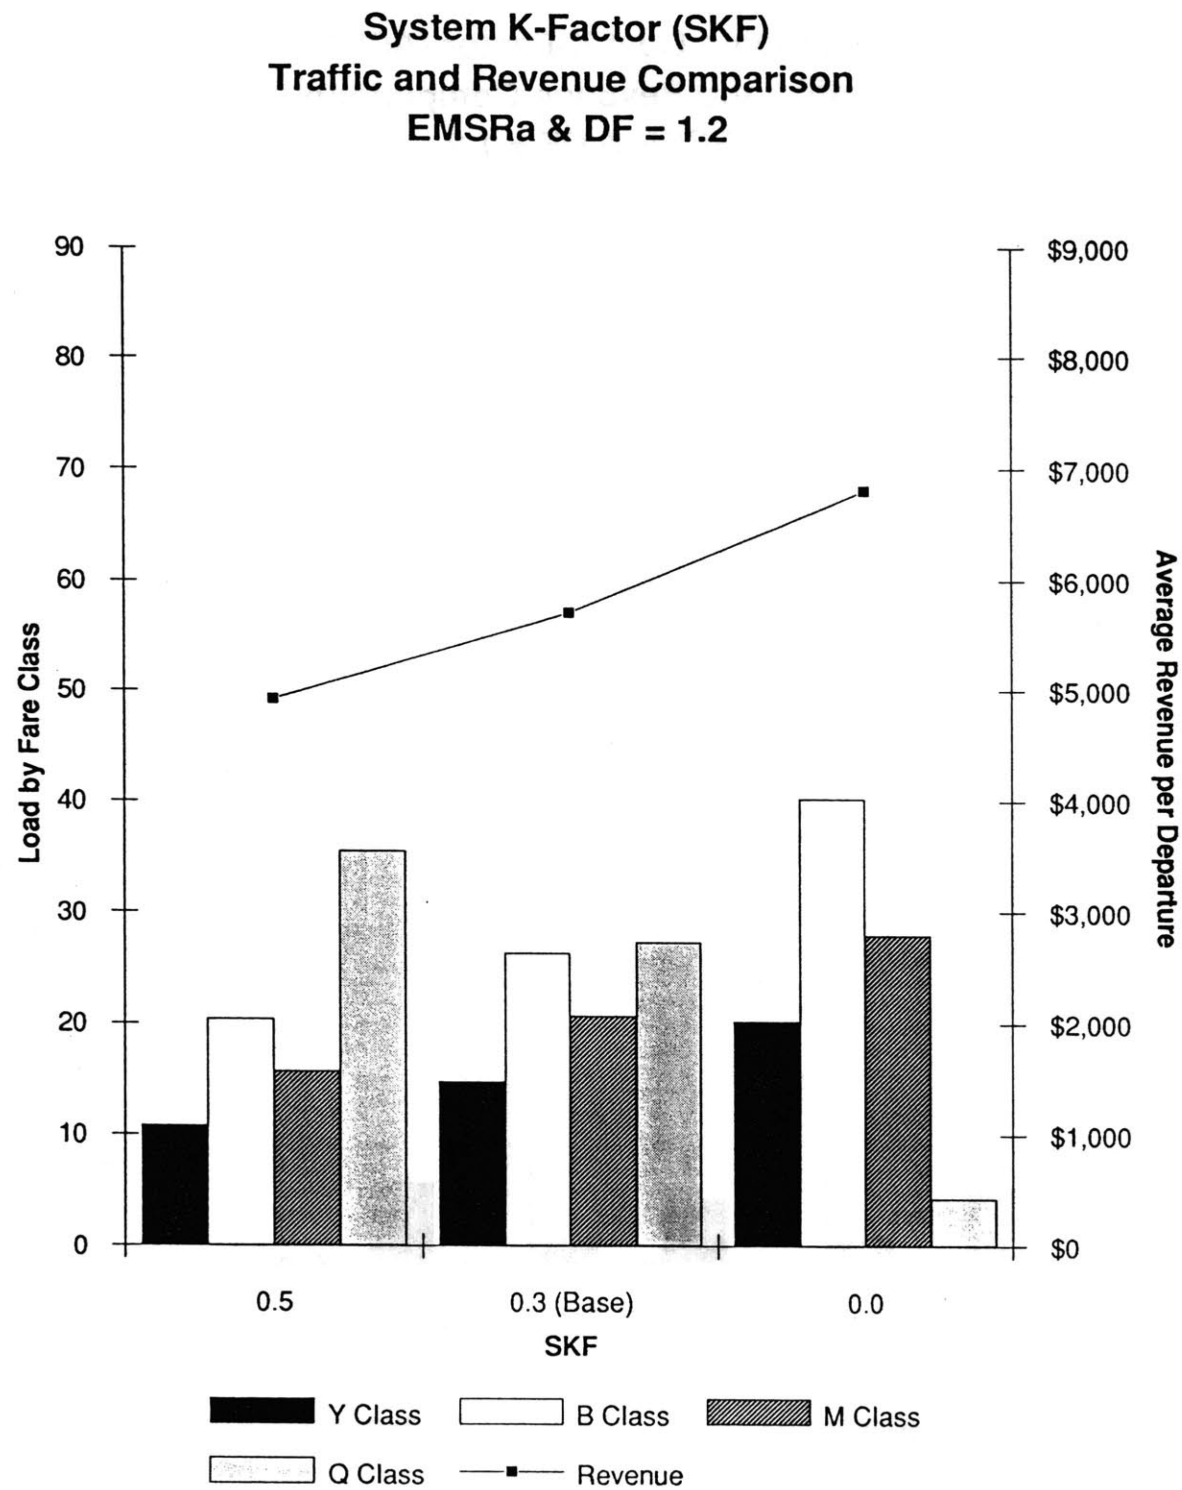

Figure 4.9. Sensitivity of System-Level Demand K-Factor, EMSR

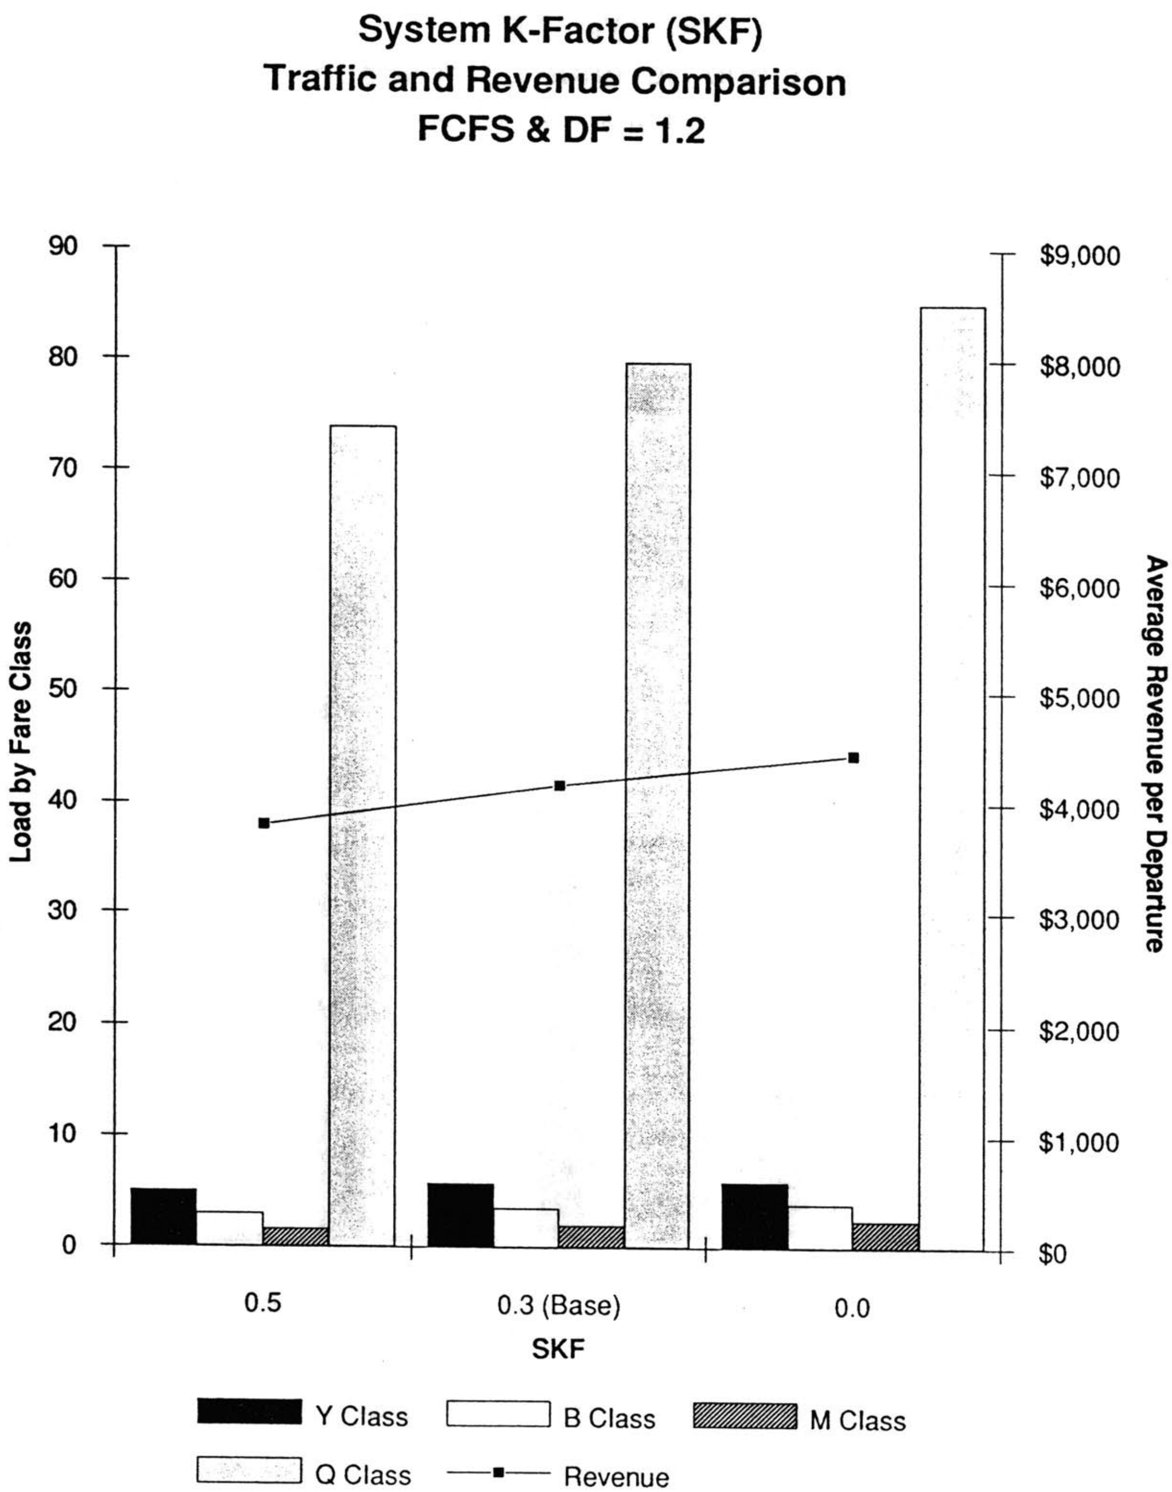

Figure 4.10. Sensitivity of System-Level Demand K-Factor, FCFS

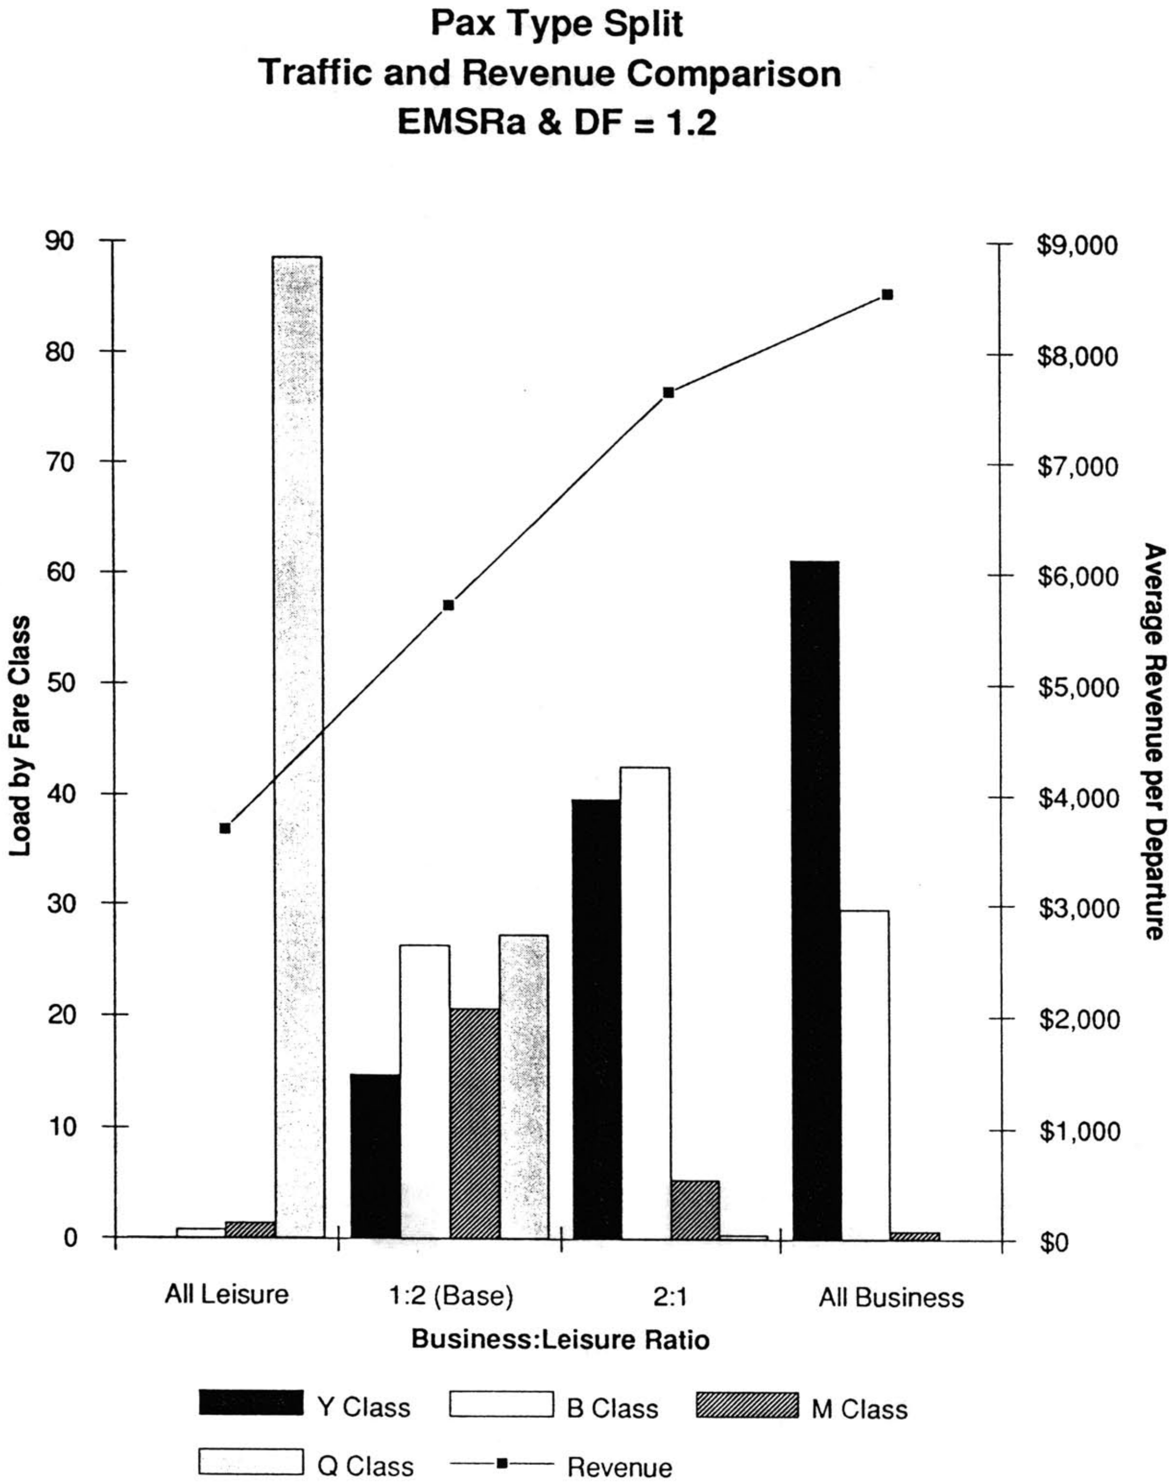

Figure 4.11. Sensitivity of Passenger Type Demand Split, EMSR

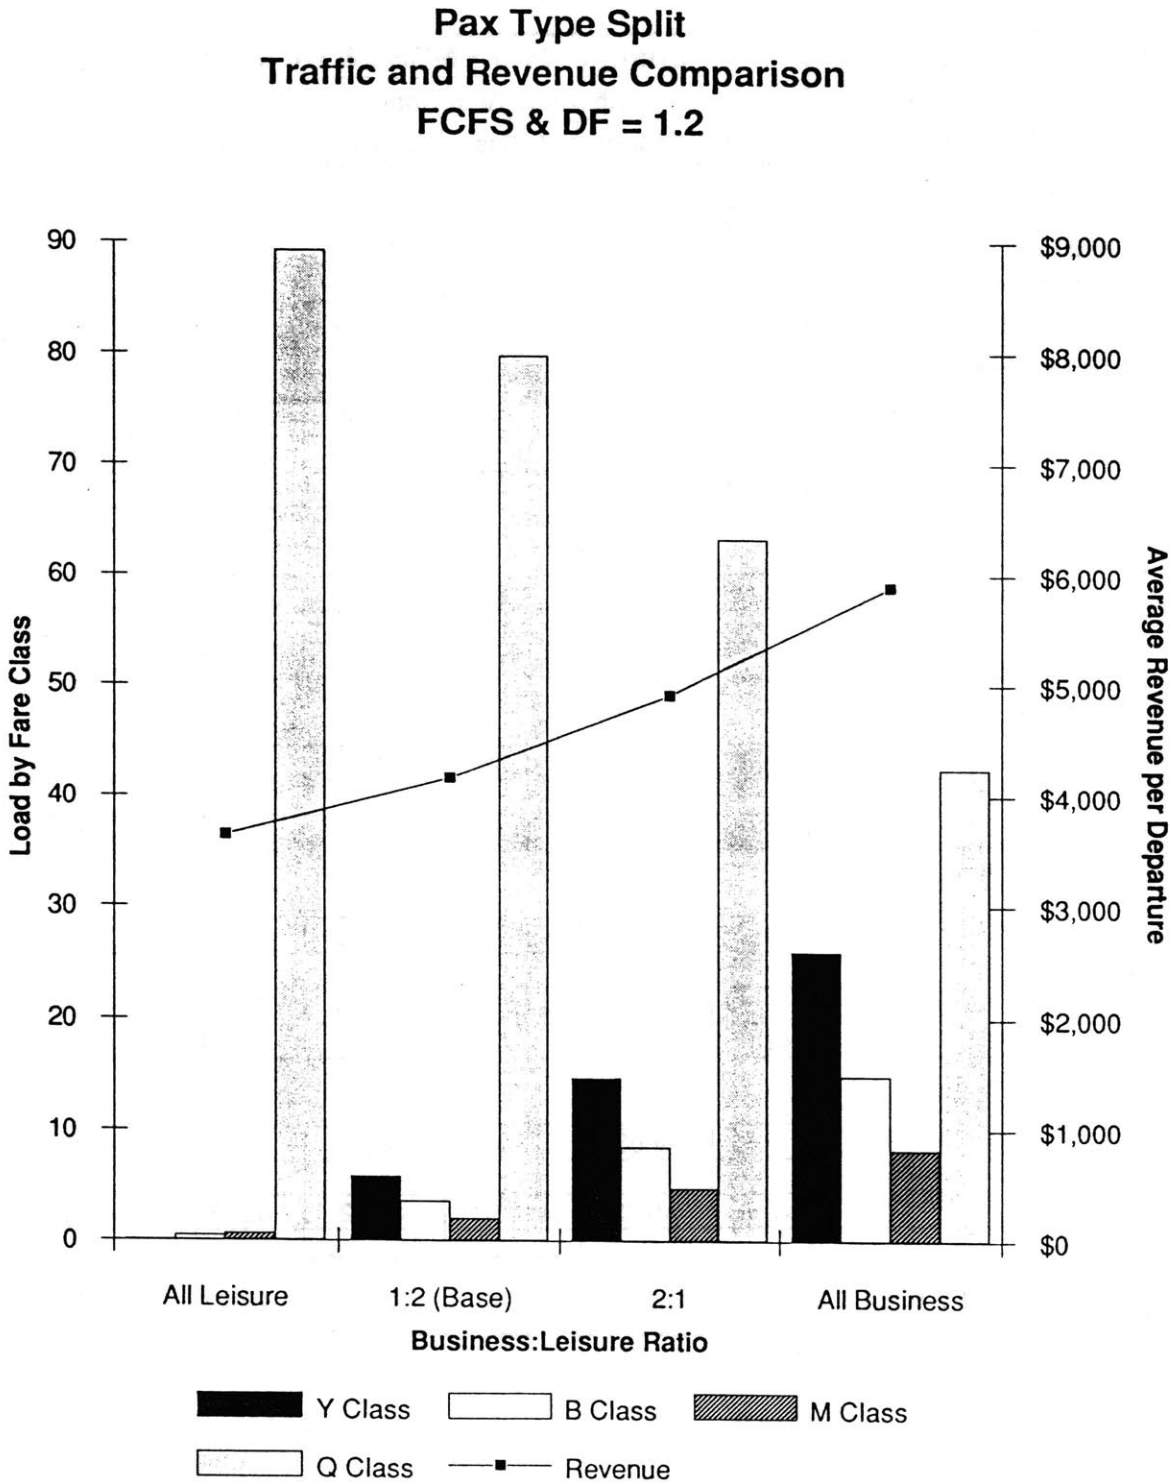

Figure 4.12. Sensitivity of Passenger Type Demand Split, FCFS

Varied Attribute and Alternate Value/Setting	Change Over EMSRa & DF = 1.2 Base					
	Absolute Seat Change					% Revenue Change
	Total	Y Class	B Class	M Class	Q Class	
Acceptable Price Ratio						
Equal to Class M Fare	-0.80	-0.10	-18.68	-13.74	+31.73	-16%
Equal to Class B Fare	-	-	-	-	-	-
Greater than Class Y Fare	+2.22	+0.67	+25.72	-15.32	-8.86	+18%
Advance Purchase						
None	+2.50	-12.22	-22.62	-19.95	+57.30	-30%
Class Q Only	+2.37	-8.03	-13.10	-0.65	+24.14	-16%
Classes M & Q	-	-	-	-	-	-
Average Restriction Cost						
Less than Class Q Fare	-0.16	-3.30	-1.89	+6.43	-1.39	-4%
Equal to Class M Fare	-	-	-	-	-	-
Equal to Class B Fare	+0.30	+3.83	+2.64	-6.95	+0.78	+5%
Attributed Cost K-Factor						
0.5	+0.22	-1.06	-2.89	+0.03	+4.14	-3%
0.3	-	-	-	-	-	-
0.0	+2.12	+9.20	+9.10	-13.75	-2.42	+15%
Passenger Type Split						
All Leisure	+1.82	-14.75	-25.48	-19.26	+61.32	-35%
1:2	-	-	-	-	-	-
2:1	-1.09	+24.81	+16.27	-15.33	-26.86	+34%
All Business	+2.65	+46.49	+3.34	-19.94	-27.24	+50%
System K-Factor						
0.5	-6.60	-4.03	-5.90	-4.95	+8.28	-14%
0.3	-	-	-	-	-	-
0.0	+3.60	+5.42	+13.91	+7.22	-22.95	+19%
EMSRa DF = 1.2 Base	89.04	14.79	26.30	20.69	27.27	$5,708

- indicates base attribute setting

Table 4.2. Sensitivity Results for EMSR Control, Demand Factor = 1.2

Varied Attribute and Alternate Value/Setting	Change Over FCFS & DF = 1.2 Base					
	Absolute Seat Change					% Revenue Change
	Total	Y Class	B Class	M Class	Q Class	
Acceptable Price Ratio						
Equal to Class M Fare	-0.97	+0.31	-0.39	-0.31	-0.58	-1%
Equal to Class B Fare	-	-	-	-	-	-
Greater than Class Y Fare	+0.08	+0.36	+0.31	+0.02	-0.60	+1%
Advance Purchase						
None	+0.57	-3.72	-1.48	-2.05	+7.82	-7%
Class Q Only	+0.40	-2.33	-0.04	+3.00	-0.22	-2%
Classes M & Q	-	-	-	-	-	-
Average Restriction Cost						
Less than Class Q Fare	-0.03	-2.15	-2.31	+1.76	+2.68	-5%
Equal to Class M Fare	-	-	-	-	-	-
Equal to Class B Fare	-0.17	+6.36	+6.37	-1.19	-11.71	+15%
Attributed Cost K-Factor						
0.5	-0.46	+0.78	+0.96	+0.04	-2.25	+2%
0.3	-	-	-	-	-	-
0.0	-0.21	+1.69	-2.80	-1.32	+2.21	-1%
Passenger Type Split						
All Leisure	-0.81	-5.75	-3.12	-1.39	+9.45	-12%
1:2	-	-	-	-	-	-
2:1	-0.07	+8.87	+4.87	+2.71	-16.51	+18%
All Business	+0.24	+20.18	+11.29	+6.15	-37.37	+42%
System K-Factor						
0.5	-7.46	-0.67	-0.53	-0.40	-5.86	-9%
0.3	-	-	-	-	-	-
0.0	+6.28	+0.16	+0.38	+0.44	+5.30	+7%
FCFS DF = 1.2 Base	91.20	5.78	3.59	2.05	79.78	$4,159

- indicates base attribute setting

Table 4.3. Sensitivity Results for FCFS Discipline, Demand Factor = 1.2

Varied Attribute and Alternate Value/Setting	Change Over EMSRa & DF = 0.9 Base					
	Absolute Seat Change					% Revenue Change
	Total	Y Class	B Class	M Class	Q Class	
Acceptable Price Ratio						
Equal to Class M Fare	-1.11	+0.45	-3.26	-4.58	+6.28	-5%
Equal to Class B Fare	-	-	-	-	-	-
Greater than Class Y Fare	+1.35	+0.55	+3.55	+0.77	-3.52	+6%
Advance Purchase						
None	+1.52	-7.64	-5.27	-7.59	+22.02	-16%
Class Q Only	+1.50	-5.04	-2.03	+0.23	+8.33	-8%
Classes M & Q	-	-	-	-	-	-
Average Restriction Cost						
Less than Class Q Fare	-0.55	-1.97	-0.20	+10.08	-8.46	-1%
Equal to Class M Fare	-	-	-	-	-	-
Equal to Class B Fare	+0.11	+5.09	+4.25	-5.39	-3.84	+10%
Attributed Cost K-Factor						
0.5	-0.12	+0.05	-0.38	-3.18	+3.38	-1%
0.3	-	-	-	-	-	-
0.0	+0.43	+5.51	+1.22	+0.71	-7.01	+10%
Passenger Type Split						
All Leisure	-0.08	-9.88	-6.86	-6.52	+23.18	-22%
1:2	-	-	-	-	-	-
2:1	+0.46	+13.39	+6.24	+1.89	-21.06	+26%
All Business	+1.32	+29.79	+13.00	-0.57	-40.90	+56%
System K-Factor						
0.5	-5.22	-1.64	+0.72	-0.16	-4.15	-7%
0.3	-	-	-	-	-	-
0.0	+5.24	+1.42	-1.49	-3.41	+8.73	+5%
EMSRa DF = 0.9 Base	79.74	9.94	7.92	7.83	54.05	$4,181

- indicates base attribute setting

Table 4.4. Sensitivity Results for EMSR Control, Demand Factor = 0.9

Varied Attribute and Alternate Value/Setting	Change Over FCFS & DF = 0.9 Base					
	Absolute Seat Change					% Revenue
	Total	Y Class	B Class	M Class	Q Class	Change
Acceptable Price Ratio						
Equal to Class M Fare	-0.89	+0.20	-0.61	-0.46	-0.03	-1%
Equal to Class B Fare	-	-	-	-	-	-
Greater than Class Y Fare	+0.85	+0.47	+0.49	-0.01	-0.10	+2%
Advance Purchase						
None	+1.15	-5.90	-2.07	-2.75	+11.88	-11%
Class Q Only	+1.41	-3.77	+0.09	+4.77	+0.32	-3%
Classes M & Q	-	-	-	-	-	-
Average Restriction Cost						
Less than Class Q Fare	+0.25	-2.31	-2.34	+2.25	+2.65	-5%
Equal to Class M Fare	-	-	-	-	-	-
Equal to Class B Fare	+0.39	+5.37	+5.57	-1.62	-8.93	+14%
Attributed Cost K-Factor						
0.5	+0.14	+0.35	+0.72	-0.02	-0.90	+1%
0.3	-	-	-	-	-	-
0.0	+0.52	+2.93	-3.11	-1.78	+2.48	+1%
Passenger Type Split						
All Leisure	-0.43	-7.95	-3.40	-1.65	+12.57	-17%
1:2	-	-	-	-	-	-
2:1	+0.70	+9.95	+4.17	+2.10	-15.52	+21%
All Business	+0.73	+21.00	+8.79	+4.09	-33.15	+43%
System K-Factor						
0.5	-5.43	-1.53	-0.73	-0.57	-2.59	-9%
0.3	-	-	-	-	-	-
0.0	+5.02	+2.02	+0.81	+0.58	+1.61	+9%
FCFS DF = 0.9 Base	80.06	8.02	4.23	2.75	65.06	$3,880

- indicates base attribute setting

Table 4.5. Sensitivity Results for FCFS Discipline, Demand Factor = 0.9

5 Competitive Scenarios

The discussion of the Operational Competitive Simulation Environment (OCSE) and the sensitivity tests of the previous two chapters have set the stage for the central mission of this thesis—the modeling of RM in competitive market settings. Using this framework, Chapter 5 directly addresses the question of whether the value of RM depends not only on the conventional demand parameters identified in previous booking optimization studies, but also on the inventory management policies of rival carriers. Because the nature of competition in a market determines the interactions between individual carriers and prospective passengers, three classes of competitive scenarios have been studied. In increasing order of complexity, these are: symmetric two path scenarios for one OD pair, dominant carrier scenarios for one OD pair, and three-city scenarios.

One subsection is devoted to each class of scenarios, containing a description of the general competitive conditions, including the rationale for the dimensions tested. Following this introduction, the results of the simulation experiments are presented in the form of graphs, tables, and interpretive text. The analysis of these results emphasizes two issues. First, it is important to understand how the existence of RM affects revenue, total enplanements, and fare class distribution in aggregate or industry-wide terms. Beyond this, the question arises of how carriers with different RM capabilities share these revenue benefits. Collectively, the answers to these questions provide evidence that the competitive circumstances in a market are indeed a determinant of the value of RM innovation.

5.1 Symmetric Two Path Scenarios

In this simplest classification of competitive scenarios, two airlines serve an isolated OD market, each with one daily departure. All competitive dimensions for the two airlines (denoted Carriers 1 and 2) are identical, with the possible exception of RM control method. In particular, the two available flights share a common departure time and flight duration. Therefore, to test the three leg-based control alternatives of FCFS, EMSRa, and EMSRb, six method *combinations* must be considered. Because RM control is the only distinguishing feature in the market, it is unnecessary to evaluate the *permutations* of the three methods, of which there are nine. For instance, the scenario, "Carrier 1 (EMSRa) versus Carrier 2 (FCFS)" produces the same results (within the range of sampling error) as "Carrier 1 (FCFS) versus Carrier 2 (EMSRa)."[36]

Sets of trials for the six combinations have been simulated at two demand factors (DFs). The same DF values used in the sensitivity analysis of Chapter 4, 0.9 and 1.2, are applied here. Previous research (Tan, 1994) indicates that the lower bound on the range of DFs relevant to RM (over which inventory control would be expected to produce a significant revenue effect) is 0.7. A slightly higher value of 0.9 has been chosen here with the purpose of highlighting any potential differences among the EMSR variants. Interestingly, simulation results which will be discussed below (see Figures 5.4-5.6) bear out this early hypothesis. The second value of 1.2 represents the upper end of the DF range which might reasonably be maintained in an actual market.

One instructive approach for interpreting the statistics for the six method combinations recalls the introduction to the thesis, which identified three phases in the emergence of RM. In the beginning, all carriers stood on an equal, unmanaged FCFS footing. At some point, one carrier—the *innovator*—would acquire an elementary RM capability in leg-based EMSR control.[37] Spurred on by the knowledge that a better "mousetrap" has been discovered, and anxious to share in the resultant revenue benefits, the *laggard* rival eventually recovers, perhaps with third-party assistance, and regains parity with the innovator. While this thought experiment is an intuitive one, no previous simulation method has sought to quantify carrier and market revenues at each stage of the history, nor to explore the changes in traffic patterns which would create these revenue effects.

Figures 5.1 through 5.3 and Table 5.1 address precisely these questions. First, to illustrate the average revenues achieved under each method combination at a DF of 0.9, Figure 5.1 graphs percentage changes over a stable base—the initial conditions in which both carriers operate under an FCFS discipline. The first two sets of columns show the situations after the innovating carrier has implemented each version of EMSR control. The percentage gain won by the innovator amounts to 8% with EMSRa. This figure is a bit higher, about 10%, with EMSRb.

---

[36]Under PODS' passenger assignment procedure, differentiating departure time alone has no effect on carrier market share. Consequently, an asymmetric flight schedule need not be simulated.
[37]Perhaps by virtue of greater investment in human capital, i.e. the recent hiring of graduates from programs in transportation economics and operations research.

Past estimates of the revenue enhancement potential for leg-based EMSR, from booking optimization models which simulate an isolated flight, have found a somewhat lower range of 5-6%. From the OCSE results, which incorporate the interactions in bookings for competing flights, it seems the presence of an FCFS rival increases the dividend from adopting EMSR. This important finding may be traced to the handling of early discount class booking requests. When EMSR control is adopted against an FCFS competitor, the uncontrolled flight functions as a "sponge," accepting all requests it receives. This has two beneficial consequences for the EMSR carrier. First, as long as one low-fare ticket is available in the market (offered on the FCFS flight), the controlling carrier will realize little demand for its discount classes. Second, by the time most high-yield business passengers appear later in the booking process, the FCFS flight will be nearly full, eventually leaving the EMSR carrier with a monopoly position. These effects are not present in an isolated market scenario—all passengers must seek accommodation on the only flight offered in the market—limiting the potential revenue gain from RM.

The new result arises with the ancillary penalty incurred by the stationary rival. That there should be a penalty at all requires some explanation. Clearly the benefit derived by the innovator stems from an improvement in the mix of bookings by fare class afforded by the control algorithm. Specifically, for flights in which leisure demand may have formerly crowded out late-arriving business passengers, surplus requests for Q class will now be turned away to reserve space for higher-yield traffic. This effect is depicted in the left-hand side of Figure 5.2, where, in the transition from FCFS/FCFS to EMSRa/FCFS, the innovator accepts fewer Q class bookings and increases loads (marginally in absolute terms but dramatically when expressed in percentages) in all other classes. Total load also drops slightly, shown by the dip in the line just below the "Innovator" label.

The penalty to "Laggard," then, occurs because those leisure passengers refused on the innovator's controlled flight still wish to travel, and will consequently seek reaccommodation on the only other available departure opportunity in the market. In the right-hand side of Figure 5.2, the extra Q class bookings appear, causing the nominal traffic in higher classes to dwindle further still. Total traffic rises by the same amount as the decline for Innovator, leaving the market passenger count unchanged.

Returning to the revenue side of the equation, Figure 5.1 reveals that Laggard's penalty is limited to less than 2% of the base FCFS/FCFS combination, regardless of the type of EMSR forged by Innovator. This liability is small when compared to Innovator's 8-10% reward. Total market revenues thus grow by 3-4%. Evidently, RM advantage does not follow the zero-sum game rules observed in contests for market share—in which gains for one participant necessarily come at the expense of a rival. The natural follow-up question to this finding looms: whence, then, does the residual revenue windfall spring?

Given the simplicity of the scenario, it is not difficult to determine that the source must be the only other actor in the marketplace—the passengers. Only part of Innovator's benefit follows from the substitution of business for leisure passengers, who are subsequently redirected to Laggard. The rest is due to the buying up behavior of those leisure passengers with a high enough acceptable price ratio, upon finding the Q or M discount classes closed. This extra out-of-pocket expense could not possibly detract from Laggard's earnings. Speaking in microeconomic terms, the innovating producer has captured more of the passengers' consumer surplus.

In the second transition, once Laggard recovers by fielding an equivalent inventory control system, one striking point emerges from Figure 5.1. In the third and fourth column groups, the revenue for Innovator is seen to be unaffected when Laggard "catches up." With the scenario balanced once again, Laggard reverses the 2% decline and climbs to the level formerly held exclusively by Innovator. Therefore, the "first mover" advantage which accrues to Innovator consists only of the inter-carrier revenue spread during the period of time in which the RM capabilities are unequal—Laggard receives the same benefit after the delay. Figure 5.2 demonstrates the related balancing of the fare class distributions, as an equilibrium is reached and one carrier can no longer unilaterally "dump" unwanted leisure passengers to a passive FCFS competitor. Individual fare class mixes moderate, and total loads return to their original levels.

As would be expected, the same sequence of events plays out at a DF of 1.2, with both the absolute and percentage effects magnified considerably. Because the revenue scales are so different, an analog of Figure 5.1 for the higher DF has not been prepared. Instead, Table 5.1 presents the data for both DFs, side by side. Although Innovator's advantage increases by more than a factor of eight for EMSRa, the accompanying damage to Laggard remains low. The most likely explanation for this result is that Laggard's FCFS performance simply cannot get very much worse after the innovation phase—at such a high DF, Q class bookings almost entirely overrun the distribution, whether or not a rival unloads additional leisure passengers as a result of a new control policy. Curiously, the distinction between EMSRa and EMSRb disappears in the second (Innovator-Laggard) stage—even carried out to tenths of a percentage point—but persists in the other dual-RM scenarios.

	% Revenue Change Over FCFS/FCFS					
	DF = 0.9			DF = 1.2		
Scenario	Carrier 1	Carrier 2	Total	Carrier 1	Carrier 2	Total
EMSRa/FCFS	+8.2%	-1.4%	+3.4%	+68.7%	-3.1%	+32.8%
EMSRb/FCFS	+10.4%	-1.4%	+4.5%	+68.7%	-3.1%	+32.8%
EMSRa/EMSRa	+8.2%	+8.4%	+8.3%	+43.1%	+43.2%	+43.2%
EMSRb/EMSRb	+10.3%	+11.0%	+10.6%	+46.5%	+47.0%	+46.8%
EMSRa/EMSRb	+7.5%	+7.4%	+7.4%	+42.3%	+44.8%	+43.6%

Table 5.1. Symmetric Scenario Revenue Effects

Figure 5.3 graphs the fare class redistribution processes at work for the higher DF. Although both the shift towards Y class and the load decrease for Innovator are far stronger than before, a similar equilibrium ensues after Laggard obtains EMSR control.

Despite marginally outperforming its more computationally intensive cousin in indirect comparisons, when EMSRb faces EMSRa directly (final columns of Figure 5.1), the aggregate outcome resembles that for double EMSRa, rather than double EMSRb. This occurs at both DFs but is particularly noticeable at a DF of 0.9. The reasons for this are not well understood but may be specific to the fare class structure and not a universal result.

To explore the distinction between EMSRa and EMSRb further, another graphical perspective has been presented in Figures 5.4-5.6. These charts show the own carrier revenue benefit achievable when switching from FCFS to each version of EMSR, given the presence of a static competitor who maintains FCFS (5.4), EMSRa (5.5), or EMSRb (5.6) control. The charts were initially created using just the two demand factors from the other experiments, but the gap between 0.9 and 1.2 was subsequently filled in by 10% increments. In addition, curiosity about the continuing trend of the revenue curves below a DF of 0.9 prompted the addition of a new DF lower bound of 0.8.

Collectively, the three graphs suggest that there is no appreciable difference in revenue potential for EMSRa and EMSRb at either end of the new DF range, regardless of the constant RM capability of the rival carrier. The explanation for this finding begins with the assertion that, as mentioned at the start of this subsection, at sufficiently low DFs, no control scheme can significantly improve revenues, simply because all booking requests will be accepted when aircraft capacity is not a binding constraint. Conversely, once a certain DF threshold is exceeded, the number of business passengers becomes so great that all rational inventory managers will refuse any lower-yield traffic—again, gradations between optimization algorithms become blurred.

However, for the more likely intermediate DF values, the superiority of EMSRb (again, conditioned on the current fare class structure model) becomes evident. Even more remarkable, the competitor's control discipline is an important determinant of this advantage. Figure 5.4, in which FCFS is the constant rival RM capability, shows that EMSRb can yield an improvement over the FCFS base nearly 10% greater than that provided by EMSRa, under certain demand factors. The point of greatest divergence appears to lie between the DFs of 1.0 and 1.1. If the static competitor runs EMSRa (Figure 5.5), the split again occurs around a DF of 1.0, although the spread shrinks to 5%. Finally, when pitted against itself (Figure 5.6), EMSRb holds a slight, but now constant, edge over EMSRa of 3-4% across all DFs.

Even in this most basic class of competitive scenarios, several significant and complex findings have emerged. The next subsection turns to asymmetric carrier rivalry, still within a single market.

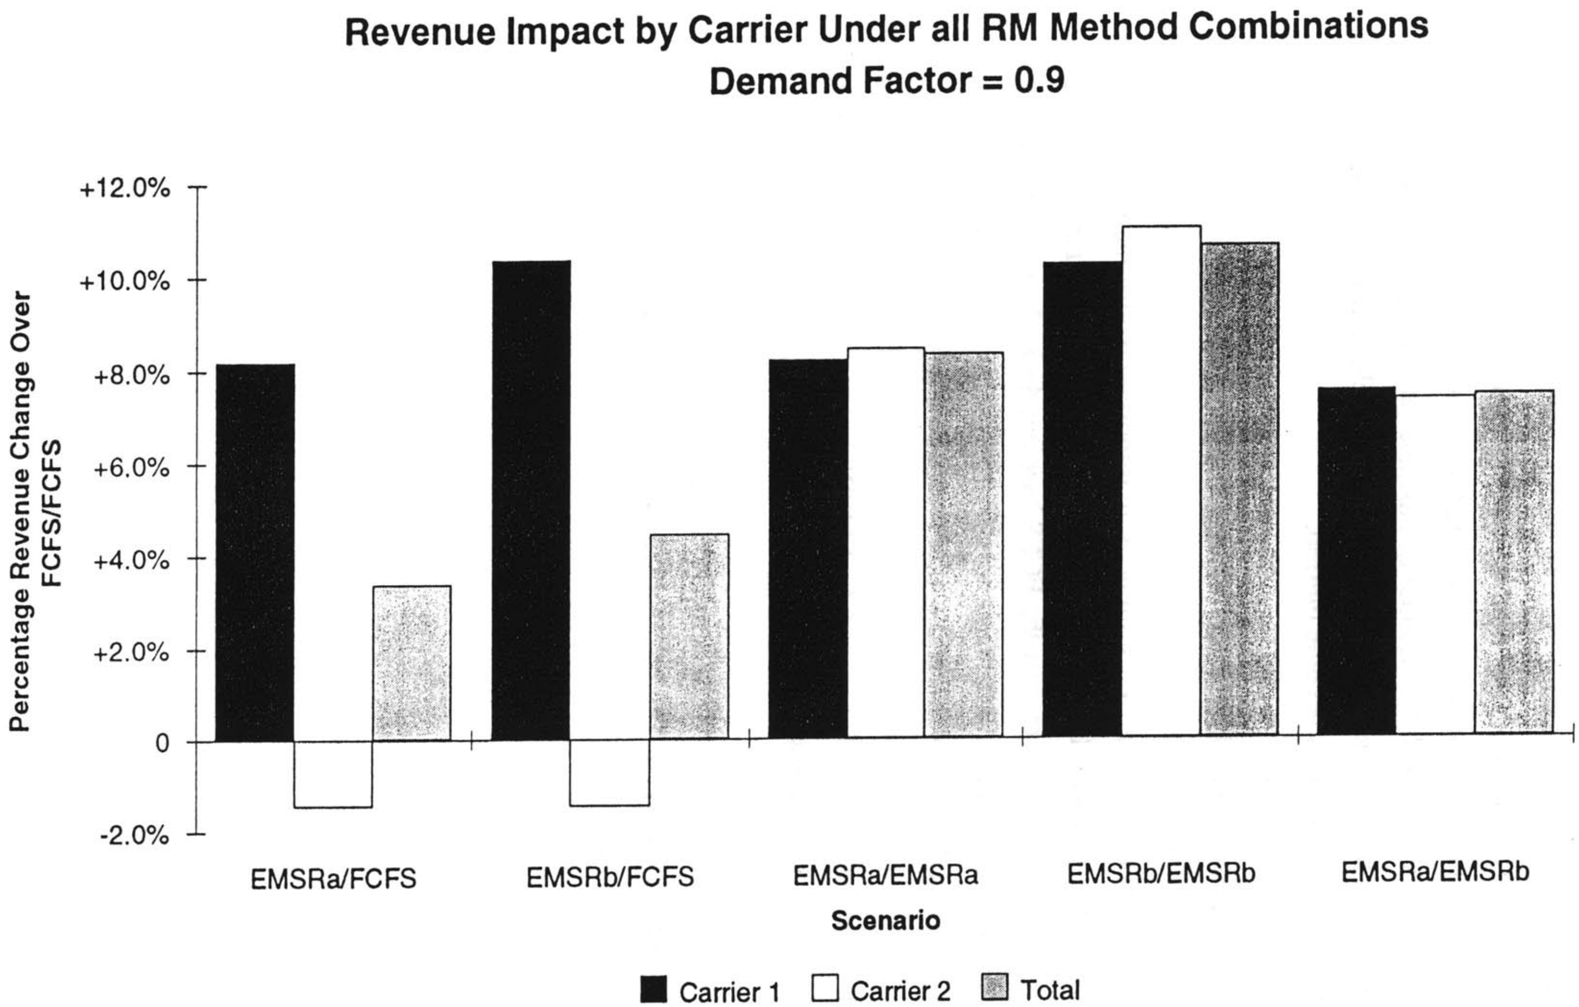

Figure 5.1. Value of Innovation, Symmetric Scenarios, Demand Factor = 0.9

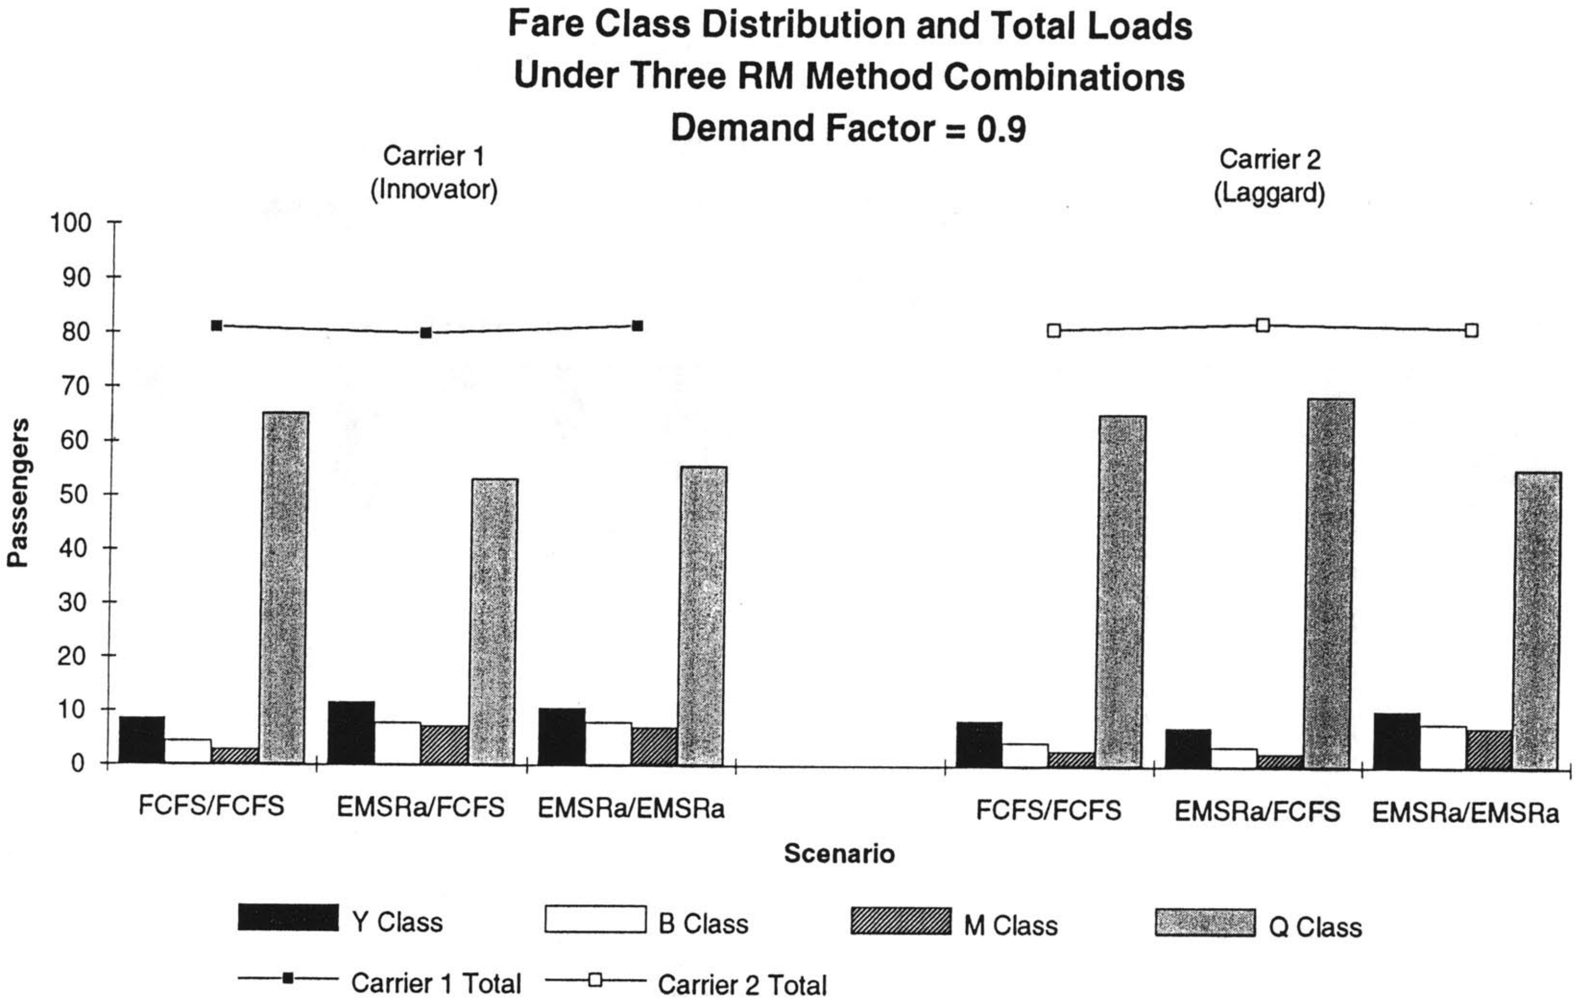

Figure 5.2. Traffic Distribution, Symmetric Scenarios, Demand Factor = 0.9

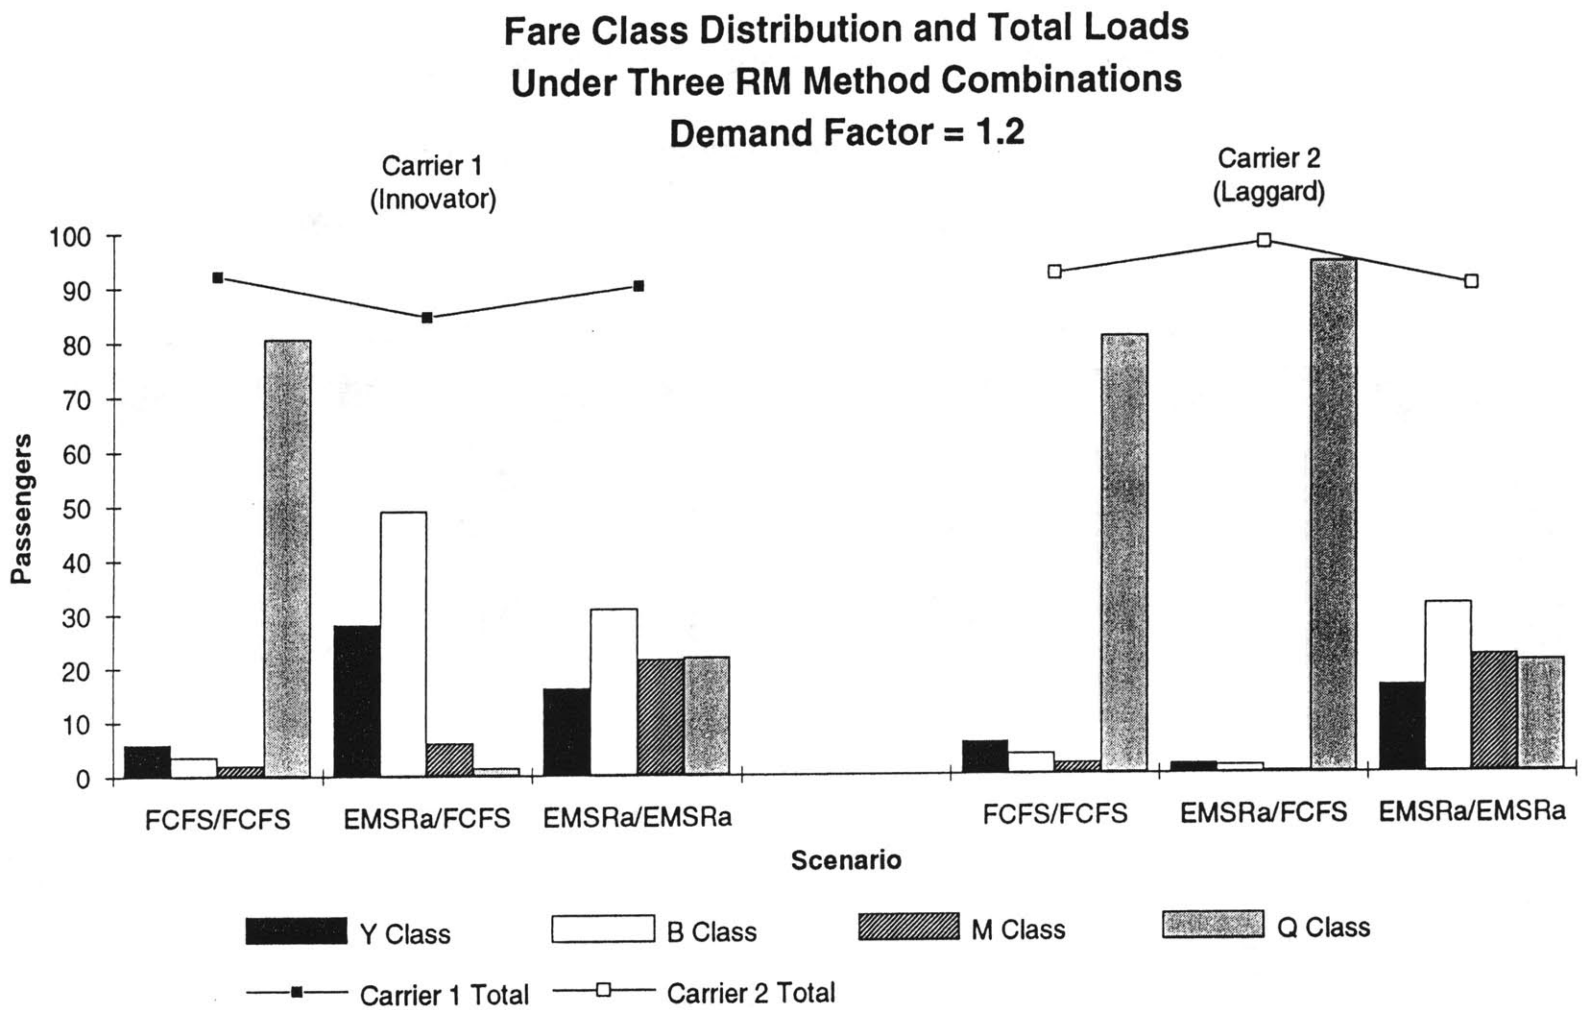

Figure 5.3. Traffic Distribution, Symmetric Scenarios, Demand Factor = 1.2

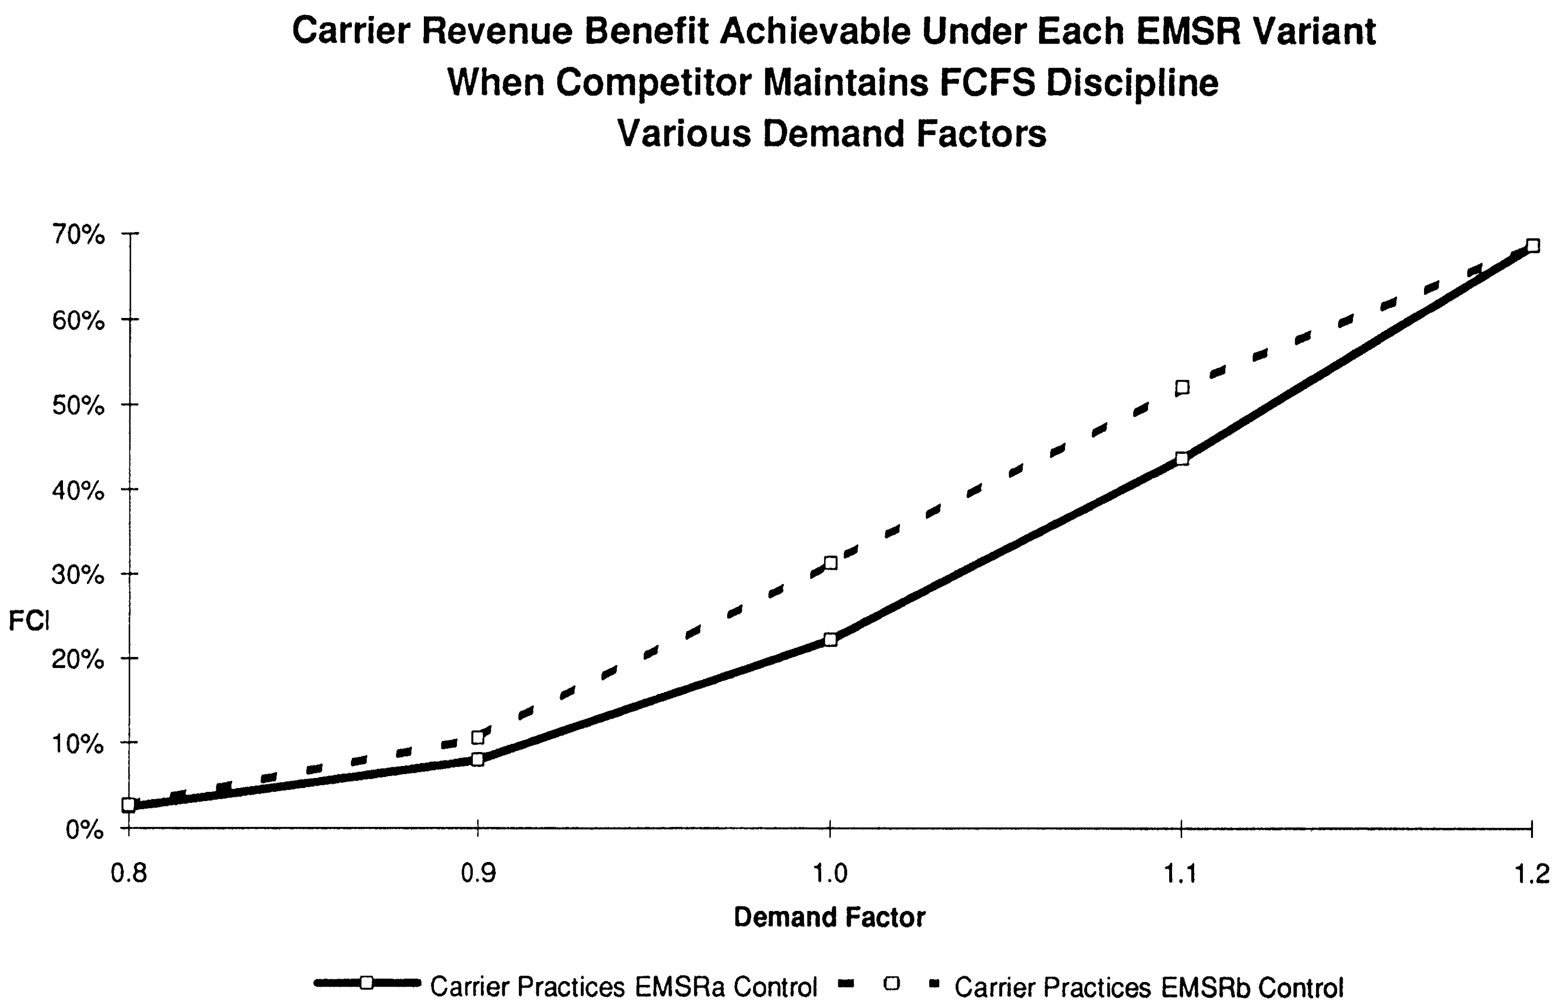

Figure 5.4. Value of EMSR Against FCFS, Various Demand Factors

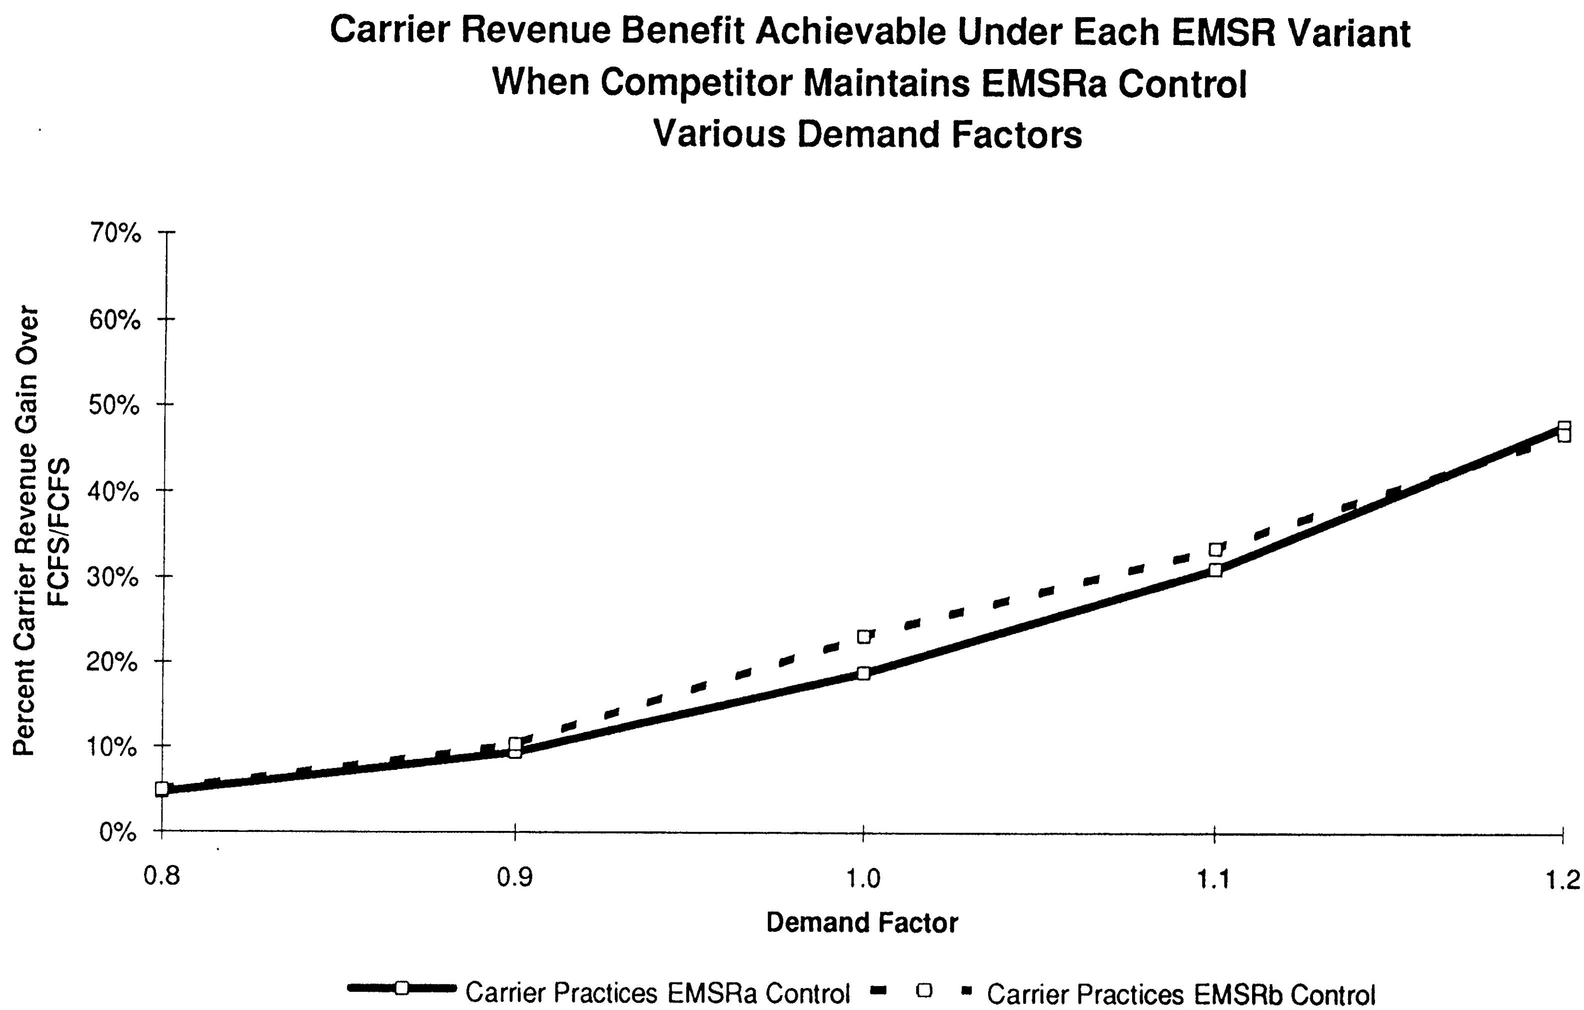

Figure 5.5. Value of EMSR Against EMSRa, Various Demand Factors

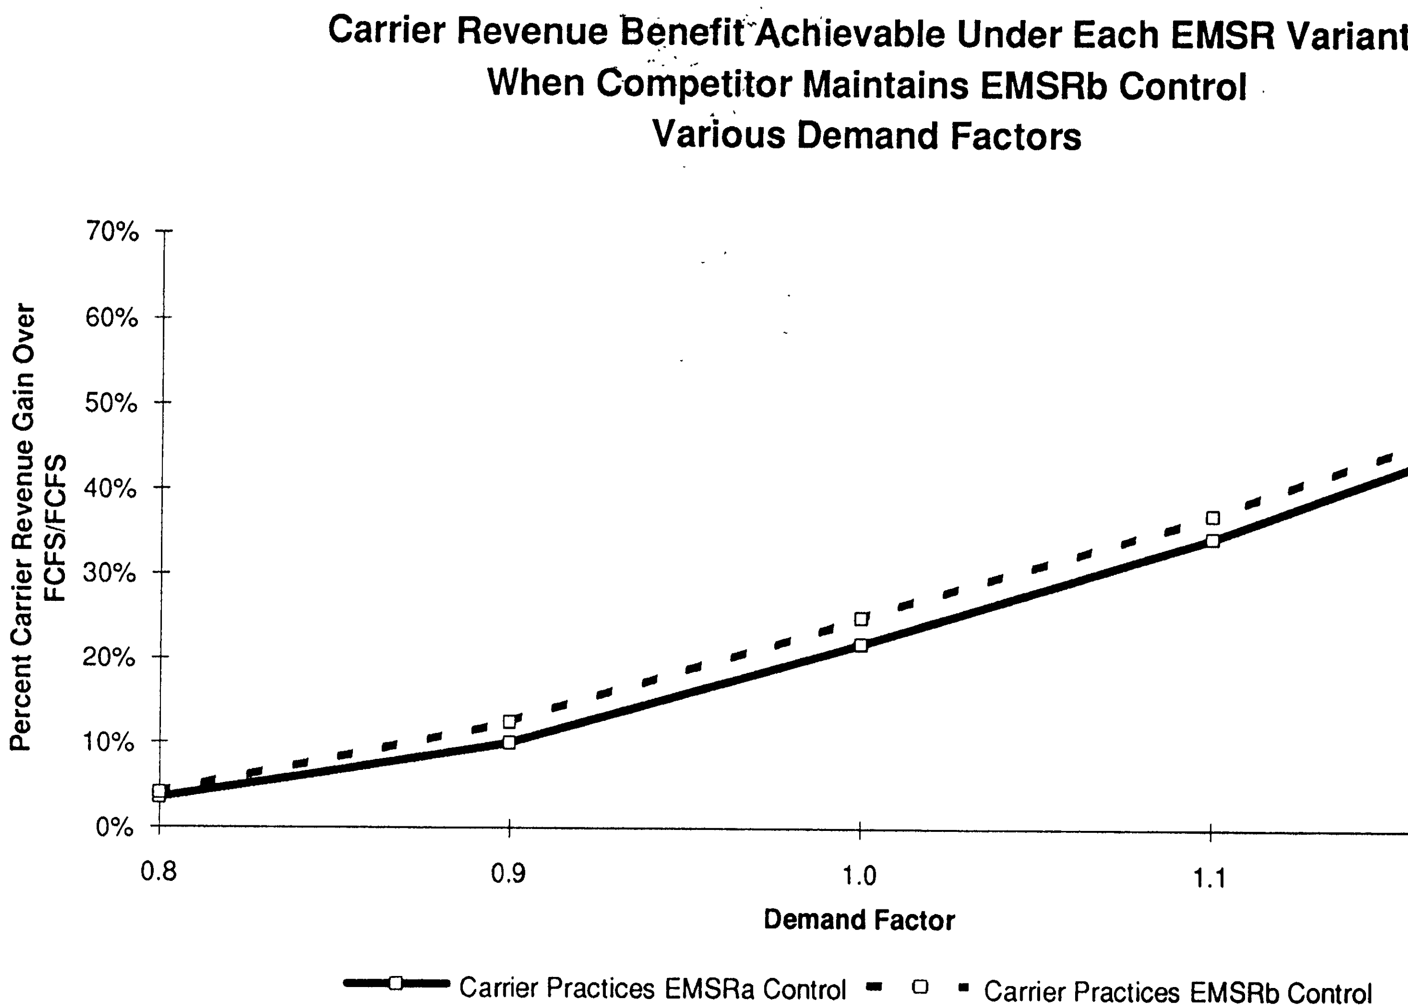

Figure 5.6. Value of EMSR Against EMSRb, Various Demand Factors

5.2 Dominant Carrier Scenarios

While the symmetric scenarios of the previous section capture many of the fundamental competitive effects found in non-monopoly markets, the nature of competition in many markets can be considerably more complex. The great majority of actual markets are not served by two identical competitors which, taken together, effectively offer just one daily departure opportunity. Typically, one carrier has a competitive advantage which causes travelers from one or more passenger segments to preferentially select certain paths—in a symmetric situation, passengers are indifferent to all available options. This competitive advantage can assume a variety of forms. One of the more common dimensions in which a carrier can distinguish its level of service is frequency. In exploring how asymmetric competitive situations alter the value of RM as demonstrated in Section 5.1, this section focuses on frequency dominance, and the associated subject of flight placement.

The natural first extension of the symmetric two path case adds one additional path to the market.[38] One carrier, the *dominant* carrier, now offers two departures to the *weak* carrier's single departure. The dominant carrier, providing twice the capacity of the weaker competitor, would expect to have higher total revenues. It is even possible that the dominant carrier will earn more on a per-flight basis. For instance, the weak carrier may receive $2,000 and the dominant $5,000, or an average of $2,500 on each of two flights. The relative revenues depend on the absolute separation of the alternate paths as well as the time-of-day placement of each departure.

To test the significance of relative flight timings, two schedules have been simulated for the dominance scenarios. In both schedules, the dominant carrier departures have been fixed at the morning and afternoon peak periods of travel, located on the default Boeing time-of-day demand distribution curves at roughly 9:00 AM and 5:00 PM. In the first, the *distinct schedule* (coded "d"), the weak carrier departure is temporally isolated from either of the strong departures. Specifically, the weak flight leaves at 1:00 PM, a point equidistant from the dominant alternatives thus maximizing the separation. Under the alternative *overlap schedule* ("o"), the weak departure coincides with the dominant afternoon peak departure.

[38]All paths in the dominant carrier scenarios continue to have the same elapsed flight times.

These extreme schedules were constructed to investigate how the separation of competing departures influences the benefits from RM. However, by holding constant the dominant path times and varying the placement of the weak departure, an implicit trade-off between frequency and time-of-day distribution has been introduced in the two schedules.

The weak carrier should have a higher revenue under the distinct schedule, all else equal, because such an arrangement gives the weak carrier a quasi-monopoly during the middle of the day. The PODS state probabilities (Section 3.3) for a daily schedule consisting of departures at 9:00 AM, 1:00 PM, and 5:00 PM show that only a small fraction of all passengers have decision windows which include two paths—and the percentage is smaller still when considering only the critical business passenger type. Therefore, many passengers will be "captive" to the weak 1:00 PM departure, in that their windows admit no other paths. In a sense, the weak carrier controls its own destiny. It does not often compete with a strong departure and so may limit the penalties (the dumping of Section 5.1) which arise with unbalanced RM capabilities in the market. However, the benefit of an exclusive market for the weak carrier is somewhat offset by its location in the off-peak. This means that if the time-of-day markets for the three paths were completely disjoint, the revenues for the weak carrier's single flight would be less than the per-flight revenue for the dominant carrier (absent RM considerations).

Alternatively, the overlap schedule places the weak departure in the afternoon peak, where the time-of-day demand distributions are most concentrated. But it is also then in direct competition with the later dominant departure. The separation of departure opportunities is even clearer under this schedule—only 2% of business passengers and 7% of leisure passengers have decision windows wide enough to initially allow them to take flights at either peak.

While it would have been desirable to isolate the frequency result from time-of-day issues, this could not be achieved without sacrificing the stable base created by the fixed dominant peak departures. However, under the Boeing curves, the confounding time-of-day effect may be minimal. In informal experiments with varying departure times, the time-of-day distributions do not fall off dramatically outside of the conventional peak areas. Instead, the distributions appear to be fairly flat with only modest peaks. Such a demand profile is in line with the long haul nature of the markets assumed in OCSE. In contrast to short haul, shuttle-type markets, which can exhibit pronounced peaks at the start and close of the business day, long haul market distributions are often more diffuse. In short, the peak / off-peak distinction is not a large one. Certainly, the frequency effect will dominate the incidental influence of time-of-day placement.

With this background, the dominant scenario results at a DF of 0.9, presented in Figure 5.7 and Table 5.2, may be interpreted. To simplify the experimental design for the more complex scenarios, only one EMSR version, EMSRa, has been selected as a representative control option.[39] Although permutations, in which the order of a pairing matters, must be considered in unbalanced scenarios, elimination of EMSRb reduces the number of permutations from nine to four. In the discussion which follows, scenarios will be described using the shorthand notation,

[Control (Strong Carrier)] / [Control (Weak Carrier)] - [Schedule Code]

For instance, "FCFS/EMSRa-d" refers to the scenario in which only the weak carrier practices inventory control under the distinct schedule, while "FCFS/EMSRa-o" represents the same control permutation but under the alternate overlap schedule.

Figure 5.7, the analog of the symmetric Figure 5.2, shows the percentage per-flight revenue changes over an FCFS/FCFS base for the three remaining permutations of control methods. The most obvious feature of the graph is the difference in average column height for the left- and right-hand sides, representing the improvement for each control permutation under the distinct and overlap schedules, respectively. In analyzing the graph, it may be most instructive to contrast the carrier revenue effects for each schedule, within a permutation, following the innovator / laggard theme from the previous section.

A convenient starting point is EMSRa/FCFS-d, because of its close resemblance to the symmetric result. The dominant innovator earns a per-flight revenue benefit of just under 10%, while the weak laggard suffers a mild loss in flight revenue of about 2%. However, market revenue rises by more than before, 5-6%, since the dominant carrier's revenues constitute a larger fraction of the market total.

In the related scenario EMSRa/FCFS-o, the outcome is greatly magnified. The dominant innovator achieves a per-flight revenue advantage of over 25%, three times the

[39]Given the findings of Section 5.1, this choice may provide conservative estimates of the value of control.

change under the distinct schedule, while the weak carrier's penalty grows to 5%. Two simultaneous forces account for this dramatic result. Both effects hinge on the same chain of reasoning. With a symmetric peak demand distribution, each peak typically receives half of the systemwide booking requests. Consequently, each peak faces an effective DF of 1.35,[40] not 0.9.

As the first consequence of this observation, because the later dominant departure and the weak departure are the closest of competitors, the former posts an extremely profitable fare class distribution. With the abundance of higher-yield traffic, the dominant flight accepts almost no Q class bookings. These refused leisure passengers quickly find accommodation on the simultaneous FCFS departure, depressing the load profile of the weak carrier.

In the morning peak, the uncontested dominant departure also fares well. Because there is no readily available FCFS alternative, and due to the extreme schedule separation, early-arrivers from the large pool of leisure passengers will receive discount class bookings. However, with the concurrently high business demand, protection levels will be set so that later leisure requests will be redirected to the afternoon departures. And, naturally, these requests will be fielded by the unmanaged flight, further worsening the weak carrier's fare class mix.

Moving to the situation in which the weak carrier is the innovator, FCFS/EMSRa-d, Figure 5.7 demonstrates that the carrier revenue changes have the same relative shapes as seen previously. Yet, both the weak carrier's gain and the dominant carrier's penalty are slightly smaller. This may be insignificant and merely a result of the incidental time-of-day complication noted above.

Unlike the first schedule contrast of the dominant innovator, FCFS/EMSRa-o shows an identical result to the distinct schedule. Naturally, the summary per-flight revenue graph masks certain flight-specific traffic shifts for both carriers. Consequently, the fare class distributions have been reviewed to determine the underlying differences in the two schedules.

---

[40]Derived by multiplying the specified flight DF, 0.9, with the number of flights, 3, and dividing by the number of effective departure opportunities, 2.

For the weak controlling carrier, total load in FCFS/EMSRa-o drops by about 6%, or five passengers, from FCFS/EMSRa-d. In revenue terms, this decrease is completely offset by an improved fare class mix. Both changes are consistent with the transition from a distinct to an overlapping flight placement. In an overlapping schedule, the weak flight has stronger competition from the dominant carrier (lower total load), but can draw away some of the high-yield business passengers in the afternoon peak (better mix). There is no net revenue change.

From the perspective of the dominant carrier, it has already been noted that the afternoon FCFS departure bears the brunt of the weak flight's relocation by accepting additional discount bookings. In order to maintain the same per-flight revenues, then, the performance for the morning departure must improve. This indeed occurs as a direct result of the schedule change—the morning peak flight, now uncontested, exhibits a better-than-average fare class mix for an unmanaged flight. Another way of interpreting the dominant carrier's constant revenue level is to imagine the peaks as the ends of a seesaw and the height of the ends as representing the revenues received by the dominant carrier on each peak flight. Initially, in FCFS/EMSRa-d, a weight, the weak departure, has been placed at the fulcrum of the seesaw, and the ends are equidistant from the ground. Continuing the analogy, under FCFS/EMSRa-o, the weight slides to one end of the seesaw, causing it to fall, while the opposite end rises. Though the distribution of heights (revenues) changes, the sum, and average, necessarily remain constant.

Finally, the third sets of columns in Figure 5.7 illustrate the final revenue positions after the laggard has recovered and implemented EMSRa. The results for scenario EMSRa/EMSRa-d mirror those for its symmetric two path counterpart. With each flight having its own distinct time-of-day "sphere of influence" in the market, and both carriers fielding equivalent control technologies, all competitors achieve an equal benefit approaching 8% of the base FCFS/FCFS-d amounts.

From the schedule contrasts of the previous two scenarios, it is apparent that the revenue pattern under an overlap schedule does not always match those obtained in a distinct schedule—the dominant carrier's advantage under EMSRa/FCFS is magnified in overlap, while the weak carrier's advantage under FCFS/EMSRa stays unchanged in overlap. Lest it be thought these are the only two possibilities for schedule differences, the dual EMSRa scenario presents still another contrast. Revenues under an overlap schedule do not even share the shape of the distinct schedule results. Where EMSR/EMSR-d shows equal unit gains for both carriers, the outcome for EMSR/EMSR-o greatly favors the dominant carrier.

To understand why this occurs requires a return to the effective DF argument presented above in the analysis of EMSRa/FCFS-o. In this dominant innovator scenario, two forces combined to produce strong gains for the dominant carrier—control on the later dominant flight produced large scale dumping of would-be Q class bookings on the coinciding weak uncontrolled flight, and the earlier uncontested dominant flight redirected many leisure passengers to the extra afternoon capacity, most of whom also subsequently found their way to the weak flight. Now, in EMSRa/EMSRa-o, the situation has changed because both afternoon departures manage their inventory, negating the first of these forces. But the second force, created by a situation in which demand is evenly divided over the peaks but capacity is not, continues to have its effect. Furthermore, this effect is clearly the stronger of the two.

With nothing distinguishing the overlapping afternoon peak flights, their fare class distributions must balance. The load statistics reveal that this is accomplished almost entirely by an upward shift of the weak departure's bookings. Whereas before the weak flight passively accepted all requests, its current control policy directs about 10%, or six, of the former Q class bookings to higher fare classes (specifically, the buy up volume is split evenly over the B and M classes). The fare mix and total load on the neighboring dominant flight remain virtually unchanged. The end result of the evolution from EMSR/FCFS-o to EMSR/EMSR-o is that the dominant carrier preserves the innovating advantage of nearly 30% while the single departure rival reverses a revenue decline and outperforms the FCFS/FCFS-o base, but only barely.

To summarize the main points of the dominance scenarios at a DF of 0.9, all results for the distinct schedule closely parallel the symmetric two path findings of Section 5.1. The distinct schedule then, can be thought of as the set of control tests against which the revenue changes of the overlap schedule may be contrasted. As anticipated, the dominant carrier's benefit from inventory control improves markedly—three times the distinct effect—under the environment of direct competition created by the overlap schedule. This occurs due to the superior performance of the uncontested dominant departure. When the dominant carrier is an RM innovator, the transferal of low-yield passengers from the later dominant flight to the overlapping unmanaged departure accentuates the revenue gap. Even when operating at an RM disadvantage (i.e., if the weak carrier innovates) the presence of an isolated peak departure limits average discount class dilution. Clearly, both RM permutation and relative flight positioning individually affect carrier and market revenues. But these experiments demonstrate that the revenue consequences of RM in a given market setting can only be evaluated by examining the interaction of these two competitive dimensions.

After the preceding analysis, which has treated in detail the dominance scenarios run at a DF of 0.9, brief mention should be made of the other tests performed in a single market, asymmetric competitive environment. Table 5.2 presents the data contained in Figure 5.7 alongside the analogous results for a DF of 1.2. Again, a graph has not been created because of the disparity in scales of the percentage revenue changes for the two DFs. Inspection of the table will reveal common trends as well as certain differences in the results for the higher DF. Application of the lines of reasoning taken above, modified for the more extreme demand conditions, resolves these disagreements.

	Per-Flight % Revenue Change Over FCFS/FCFS					
	DF = 0.9			DF = 1.2		
Scenario	Strong Carrier	Weak Carrier	Total Market	Strong Carrier	Weak Carrier	Total Market
Distinct Schedule						
EMSRa/FCFS	+9.3%	-2.2%	+5.5%	+72.1%	-3.4%	+46.9%
FCFS/EMSRa	-0.6%	+8.5%	+2.4%	-1.8%	+71.2%	+22.5%
EMSRa/EMSRa	+7.6%	+7.4%	+7.5%	+49.8%	+49.3%	+49.7%
Overlap Schedule						
EMSRa/FCFS	+26.8%	-4.8%	+16.6%	+69.4%	-4.7%	+44.7%
FCFS/EMSRa	-0.9%	+8.9%	+2.3%	-2.0%	+71.8%	+22.6%
EMSRa/EMSRa	+26.9%	+2.2%	+18.9%	+57.3%	+30.1%	+48.2%

Table 5.2. Dominant Scenario Revenue Effects

For instance, the magnification of the EMSRa/FCFS-d result in EMSRa/FCFS-o discovered at a DF of 0.9 does not arise with the higher DF because, under the distinct schedule, the benefit to the dominant innovator is already so great (the fare class mix is so heavily weighted towards the full-fare classes), that repositioning of the weak departure can yield no further improvement. In fact, there is a slight decrease, as the worsened performance of the overlapping departure detracts from the per-flight revenue statistic.

However, as with a DF of 0.9, the weak carrier still suffers from having to compete directly with a dominant flight.

As a final extension of the asymmetric scenarios, the degree of frequency superiority has been varied. In these experiments, the stronger, *super-dominant* carrier has four daily departures to the weak carrier's one. The motivation for this study began with the expectation that, as market frequency (and the number of effective departure opportunities) increases, the "distinct" schedule will become less distinct. In other words, a significant number of passengers will have multiple paths within their original decision windows. As a result, the results under the overlap schedule may be closer to those for the distinct schedule.

In a sense, the super-dominant scenarios explore what might be termed the concept of "frequency saturation"[41] in an RM context—the *a priori* hypothesis being that the per-flight revenue benefit for a dominant carrier should eventually deteriorate as the frequency imbalance grows large, regardless of the positioning of the weak carrier's departure(s). If the value of RM innovation derives from the ability to divert low-yield passengers to alternate paths, this mission inevitably begins to fail as more of these alternate paths are also served by the innovating carrier. Although this super-dominant market schedule may not in fact constitute a saturated schedule from the RM perspective, it is at least a step closer than the two versus one dominance setting. Comparison of the revenue results for the two degrees of dominance may thus provide preliminary evidence for the saturation hypothesis.

Figure 5.8 and Table 5.3 present such supporting evidence, at least at the moderate DF of 0.9—more lopsided frequency shares may be required to demonstrate the hypothesis at the higher DF. The disparity between the column heights of the distinct and overlap schedules in Figure 5.7 has substantially diminished. Comparing Tables 5.2 and 5.3, the dominant carrier's innovation advantage under EMSRa/FCFS-o and super-dominance has fallen from the dominance level of 26.8% (three times the EMSRa/FCFS-o value of 9.3%) to just 14.2%.

[41]Frequency saturation traditionally refers to the the diminishing marginal profitability of adding another flight to a market characterized by a limited demand elasticity with respect to frequency. Intuitively, trip generation factors intrinsic to a market limit the number of flights which may be profitably flown in the market—beyond this limit, addition of new capacity does not stimulate new demand but only cannibalizes the existing service.

Scenario	Per-Flight % Revenue Change Over FCFS/FCFS					
	DF = 0.9			DF = 1.2		
	Strong Carrier	Weak Carrier	Total Market	Strong Carrier	Weak Carrier	Total Market
Distinct Schedule						
EMSRa/FCFS	+9.3%	-1.4%	+7.1%	+68.0%	-3.6%	+53.7%
FCFS/EMSRa	-0.6%	+12.1%	+1.9%	-1.6%	+66.8%	+12.0%
EMSRa/EMSRa	+7.3%	+10.4%	+8.0%	+53.0%	+58.6%	+54.2%
Overlap Schedule						
EMSRa/FCFS	+14.2%	-3.9%	+10.7%	+68.7%	-3.6%	+54.4%
FCFS/EMSRa	-1.0%	+6.9%	+0.6%	-1.6%	+65.6%	+11.7%
EMSRa/EMSRa	+13.1%	+3.1%	+11.2%	+56.6%	+41.2%	+53.5%

Table 5.3. Super-Dominant Scenario Revenue Effects

The contrast between the revenue results for the dominant and super-dominant scenarios demonstrates that the degree of *coverage* in a market partially determines the value of introducing RM. Coverage can be loosely defined as the level of service provided by all carriers in a market over the course of the operating day. When there are sizable gaps in market coverage, as in the overlap dominant scenarios, many passengers must replan to conform to one of the few departure opportunities, creating high effective DFs for these flights. Differences in availability of adjacent (and coinciding) flights, caused by different control capabilities, can then lead to highly unbalanced traffic distributions and revenues for individual flights. When additional service is added, filling these gaps, more passengers find a flight convenient to their initial travel plans. As a result, even uncontrolled flights have a stable passenger base by virtue of their temporal location, lowering the potential per-flight revenue advantage from RM. Further conclusions must pend the simulation of scenarios having more complex coverage profiles.

While future research may include the detailed modeling of actual competitive situations, this concludes the assessment of abstract single market scenarios for the purposes of this thesis. The final section of this chapter expands the horizons of the current isolated market experiments to an elementary multiple city network.

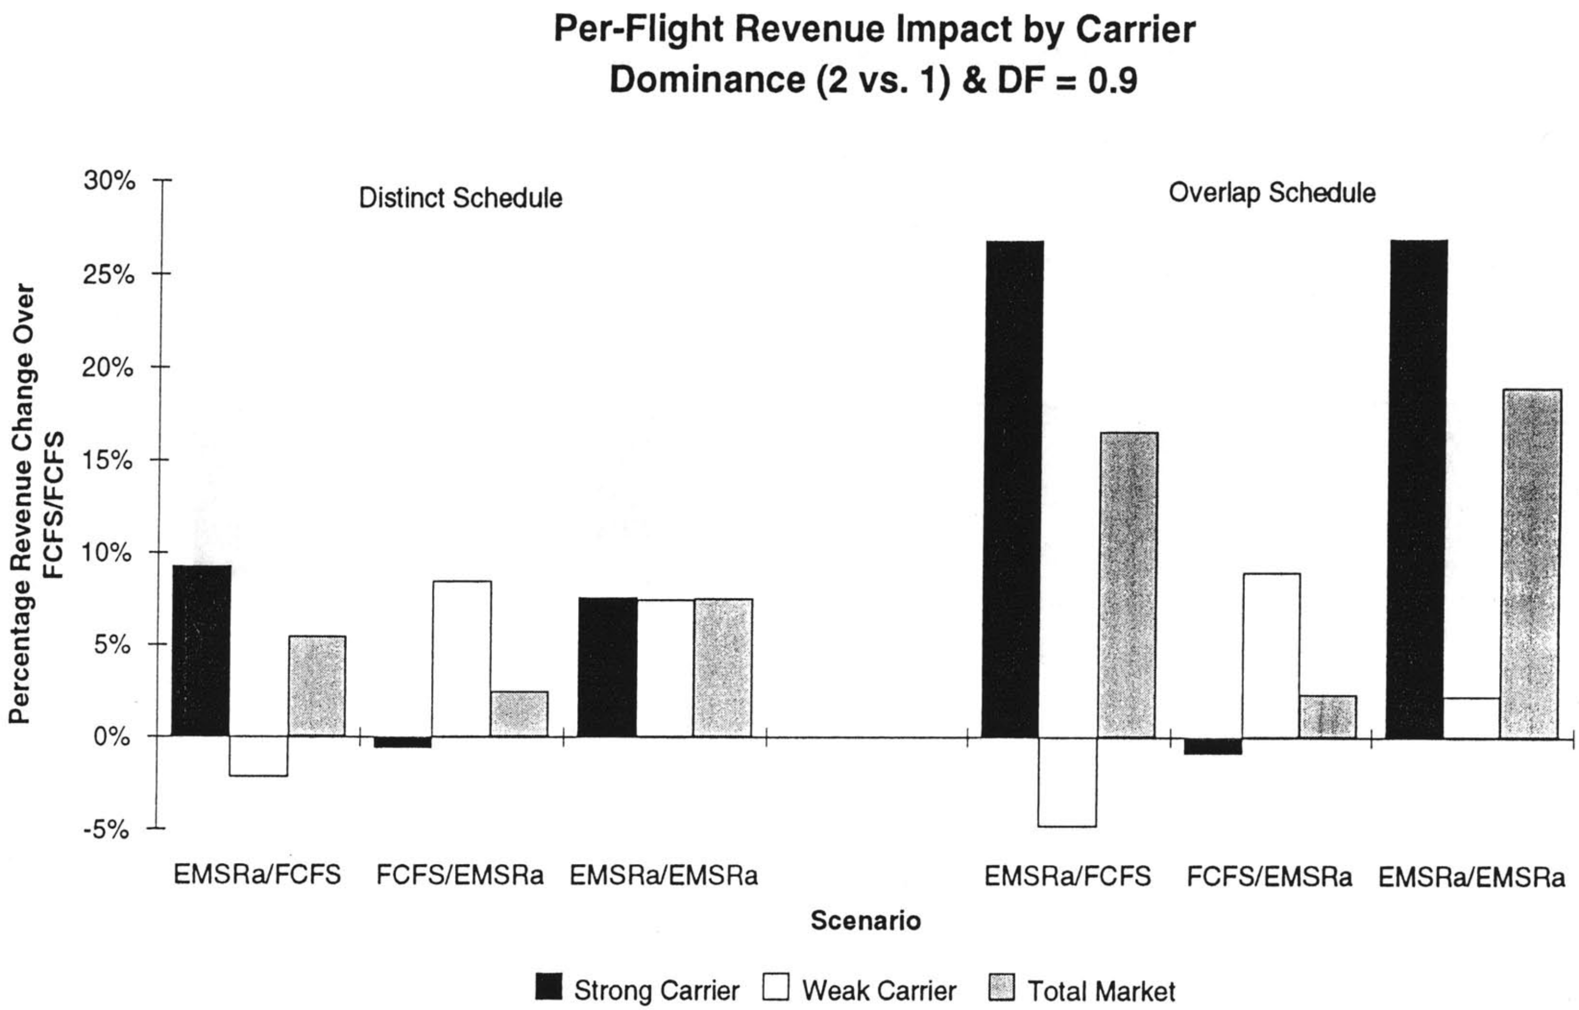

Figure 5.7. Value of Innovation, Dominant Scenarios, Demand Factor = 0.9

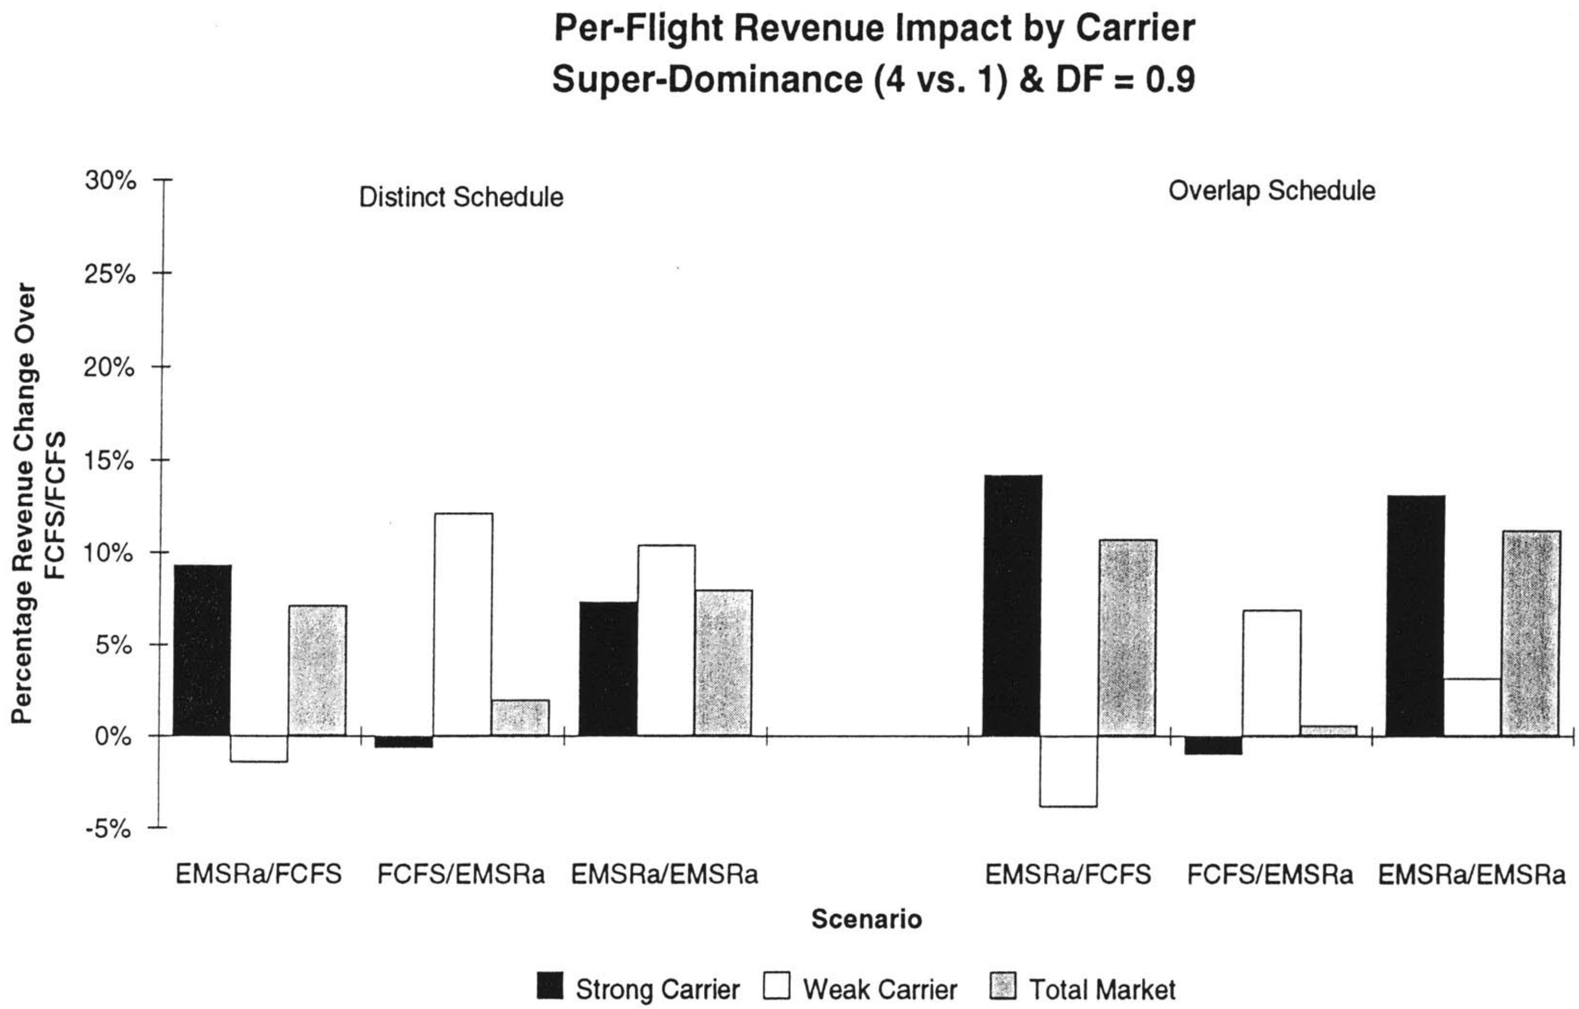

Figure 5.8. Value of Innovation, Super-Dominant Scenarios, Demand Factor = 0.9

5.3 Three-Market Scenarios

As a competitive simulation environment, even a first survey of OCSE must include results for a network containing multiple interdependent OD markets. For one reason, up to this point, one of the main features of the PODS passenger assignment framework has gone untested—the evaluation of path quality. Whether a path is nonstop, direct / through / continuing (intermediate stop with no change of plane), or connecting (intermediate stop with a change of plane) clearly affects path preference and the importance of the path quality dimension likely varies by passenger type.

The influence of path quality can be separated into two components. The stronger of these components concerns the extra time associated with paths of inferior path quality. Through and connecting paths generally have longer elapsed times, due to the layover at the intermediate station, ranging from 30 to 90 minutes at most domestic hubs, and the circuity of the path relative to the straight line vector of a nonstop departure. Longer paths are less likely to fall entirely within a randomly generated passenger decision window. Consequently, schedule states including these longer paths are assigned smaller probabilities. The severity of the reduction depends on the schedule sensitivity of a passenger type, which in turn derives from an implicit valuation-of-time, a parameter commonly found in discrete choice models. Broadly speaking, tighter schedules and higher incomes cause business passengers as a group to exhibit a high value-of-time. As a result, when deciding between two otherwise identical available paths, the connecting option will be less attractive for any passenger but particularly undesirable for the business traveler.

Furthermore, there is an intrinsic disutility to accepting a connecting itinerary, apart from the additional travel time involved. The hassle of changing aircraft and the increased possibility of misdirected luggage go into this attributed cost. The more dire risk of missing the second leg entirely, and the havoc such an eventuality would wreak on a set schedule, also must be considered. Though difficult to estimate, PODS allows the input of average attributed costs for this disutility (to which the common attributed cost k-factor is applied) and the values chosen in OCSE are \$15 and \$10 per unit of the path quality index[42] for business and leisure passengers, respectively.

---

[42]The PODS path quality index (PQI) quantifies path quality with the expression,

$$PQI = 1 + (\#stops) + 2(\#connections)$$

Besides providing an opportunity to explore another dimension of path preference, a network with overlapping markets challenges the class of RM control methods considered in the thesis. When a single leg carries traffic from more than one OD market, passengers must compete for limited capacity. From the carrier's perspective, it would be beneficial to be able to discriminate between continuing / connecting and local requests because the former typically contribute more to system revenues. It is the inability of leg-based control methods to make this distinction that has inspired research into virtual nesting and ODF alternatives. As a result of this limitation, adoption of EMSR leg-based control can have a different effect in each market served by the carrier, as will be demonstrated below.

To perform the desired experiments in path quality, and to incorporate competition for capacity in multi-path legs, the single market network of the previous scenarios must be augmented by (at a minimum) just one city. Continuing to restrict the analysis to one direction of travel, if the three cities are labeled A, B, and C, the three associated OD markets are A-B, B-C, and A-C. One OD pair, A-C, has been chosen as the "connecting,"[43] longer haul market. The remaining two local, or spoke, markets, then receive only nonstop service. Since network contribution typically differs for local and connecting markets, local fares are set at the original OCSE levels (Table 4.1) and A-C fares are uniformly 50% higher.

Two carriers participate in the network, supplying the schedules illustrated in Figure 5.9. One, the "nonstop carrier," has a single daily nonstop departure in the A-C market. The nonstop carrier (NSC) does not serve either local market. The competing carrier, the "local service carrier," offers one nonstop departure in each local market for a total of two daily flights. These local flights are timed to create a potential connecting opportunity in the A-C market. Thus, while A-C passengers may construct a connecting path from two local legs, the local service carrier (LSC) does not provide nonstop A-C service. The designations, NSC and LSC, refer to the level of service provided in the A-C market—LSC is also a "nonstop carrier" but carries A-C passengers only through its local service flights. To summarize, NSC serves only one market, A-C, with a nonstop

[43]A-C is the only market with connecting service, but because one goal of the three-city scenarios is to determine how traffic balances over paths of different path qualities, A-C service is not *exclusively* connecting. To avoid confusion, this longer haul market will be referred to as "A-C" rather than "the connecting market."

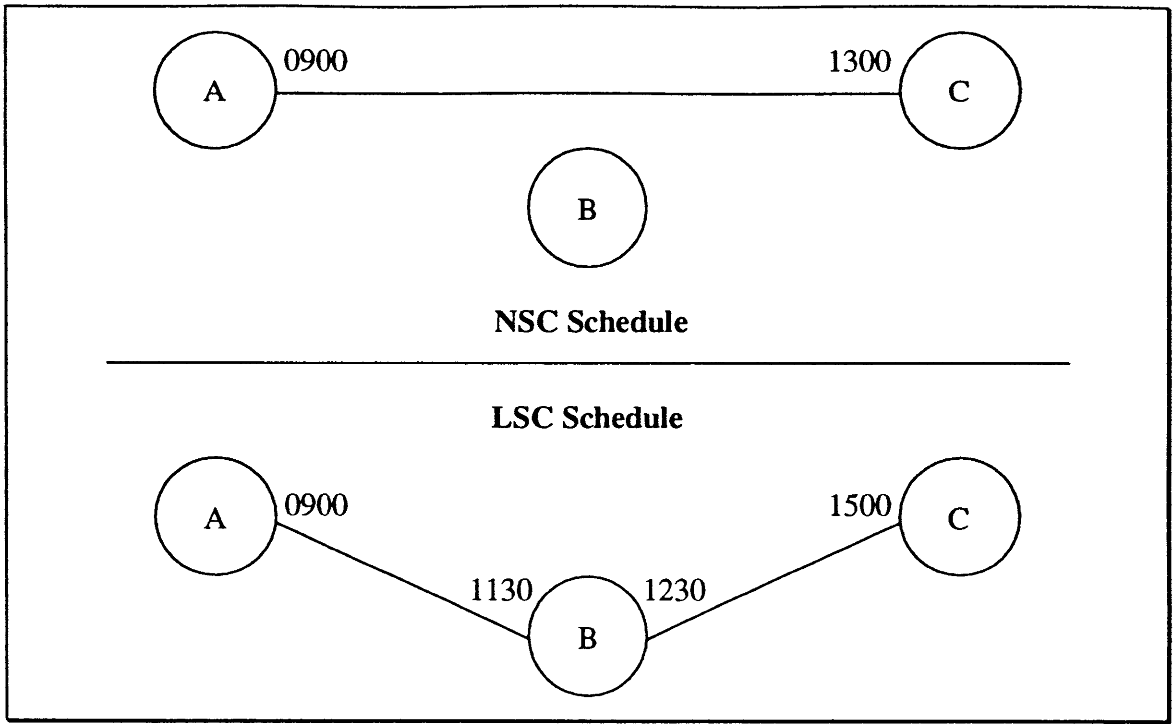

Figure 5.9. Service Schedules for Three-Market Scenarios

departure. LSC serves all three markets, with nonstop service in the two local markets and connecting service in A-C.

Within this common network structure, the distribution of system demand over the three markets may take one of three patterns. The patterns have been designed to investigate how the ratio of A-C passengers to local market passengers affects the revenue and traffic results of each carrier. This ratio is important to both carriers. Clearly, it directly influences the outcome of the competition for seats on the local legs of LSC. Fielding and accepting requests from passengers traveling in the local markets leaves less room for those attempting to connect through B in A-C. Therefore, when the system demand distribution is weighted towards the local markets, LSC revenues are low. However, the ratio also decides the success of NSC. In offering dedicated A-C service, NSC does not merely suffer depressed revenues if local market traffic predominates, but may fail to achieve adequate load factors. Conversely, NSC can take better advantage of high demand conditions in the A-C market, with service of superior path quality.

The base demand distribution is an equal split (coded "e") among the three markets. The second option, *high local* split ("h") assigns 40% of the system demand to each local market, and the remaining 20% to A-C. At the other end of the spectrum, *low* *local* split ("l") gives half of the system demand to A-C and one quarter to each local market. In all distributions, the input mean system demand is calculated as the product of network capacity (300 seats) and DF.

Before reviewing the results, the shorthand scenario notation of Section 5.2 must be adapted to the new network characteristics. The revised syntax is,

[Control (NSC)] / [Control (LSC)] - [Demand Distribution Code]

For example, in the scenario, "EMSRa/FCFS-l," NSC controls inventory, LSC does not, and the carriers face a low local market split.

PODS permits the disaggregation of RM revenue benefit by market, or class of market, for multi-city networks and this perspective is a good introduction to the analysis of the three-city scenario results. Table 5.4 offers such a disaggregation by separating the revenue effects for the A-C and local markets in each scenario. One clear finding from the table is that the revenues earned in the local markets, by LSC, show a greater percentage increase from the implementation of inventory control, regardless of the demand distribution. This follows because LSC serves multiple OD markets and thus experiences strong competition for its local capacity under any split of system demand. Within a control scenario, the rankings of revenue change by demand distribution follow intuition, with equal split always falling in the middle. Revenue from the A-C market improves the most under a low local split and the least under a high local split. The situation is reversed for the local markets. The only scenarios exhibiting significantly negative results are the EMSRa/FCFS-e and -l local market cases, in which the A-C passengers refused on NSC, and subsequently accommodated on LSC, crowd out local traffic.

Another way to view the results focuses on the changes in system (across all markets) carrier revenues. Table 5.5 indicates that the general revenue trends observed in the single market competitive scenarios also hold here, within each control permutation. However, as suggested in the analysis of Table 5.4 above, the own carrier system revenue gains in the instances when LSC innovates are magnified because of the additional competition for capacity. At an even split, LSC has an 18.1% advantage over the FCFS/FCFS base, and this benefit grows to 30% under the most favorable high local split conditions, overwhelming all other effects at this DF. Even for a low local split, the LSC

Demand/Control Scenario	% Market Revenue Change Over FCFS/FCFS			
	DF = 0.9		DF = 1.2	
	A-C Market	Local Markets	A-C Market	Local Markets
LSC Innovates				
Equal Split	+1.5%	+19.2%	+21.1%	+70.5%
High Local Split	+1.4%	+31.7%	+33.2%	+86.4%
Low Local Split	+3.6%	+9.6%	+23.7%	+50.2%
NSC Innovates				
Equal Split	+4.8%	-2.2%	+33.0%	-2.1%
High Local Split	-0.8%	+0.1%	+32.5%	+0.5%
Low Local Split	+15.6%	-6.2%	+58.0%	-20.1%
LSC & NSC Control				
Equal Split	+8.3%	+20.0%	+35.7%	+72.7%
High Local Split	+2.1%	+31.4%	+35.9%	+88.1%
Low Local Split	+18.1%	+9.9%	+76.7%	+52.8%

LSC = local service carrier
NSC = nonstop carrier

Table 5.4. Three-Market Scenario Revenue Effects by Market

Demand/Control Scenario	% Carrier System Revenue Change Over FCFS/FCFS			
	DF = 0.9		DF = 1.2	
	NSC	LSC	NSC	LSC
LSC Innovates				
Equal Split	+0.4%	+18.1%	-0.4%	+70.2%
High Local Split	+1.6%	+30.0%	+4.0%	+81.0%
Low Local Split	-0.7%	+11.5%	-0.7%	+62.6%
NSC Innovates				
Equal Split	+3.4%	-0.5%	+16.7%	+1.1%
High Local Split	-0.8%	+0.1%	+1.4%	+0.7%
Low Local Split	+16.7%	-1.9%	+57.6%	-2.4%
LSC & NSC Control				
Equal Split	+6.4%	+20.0%	+15.7%	+71.1%
High Local Split	+2.6%	+29.7%	+5.5%	+83.0%
Low Local Split	+17.6%	+12.3%	+76.4%	+57.4%

LSC = local service carrier
NSC = nonstop carrier

Table 5.5. Three-Market Scenario Revenue Effects by Carrier, Systemwide

system revenue increase of 11.5% is slightly higher than the 8-10% range discovered in the single market scenarios.

Conversely, NSC receives a smaller benefit from innovation under an equal demand split. At a high local split, there is insufficient A-C OD traffic to warrant NSC's nonstop service—load factors are so low that RM fails to generate any revenue change (the insignificant negative result is an artifact of stochastic demand variations). Only under a low local split does NSC achieve a revenue reward comparable to LSC's improvements from being the first to implement inventory control.

Table 5.5 also contains percentage data at the higher DF of 1.2. There are no results here which cannot be readily explained by the more extreme demand conditions.

However, one further perspective deserves attention. Table 5.6 disaggregates the carrier system revenues of Table 5.5 to isolate the change in LSC revenue within A-C (for NSC, serving only A-C, system and A-C revenues are equal). The remainder of this section focuses on the unusual carrier revenue effects discovered in the one market served by both competitors, A-C.

Figure 5.10 continues the series begun with the single market scenarios. Percentage revenue changes over the FCFS/FCFS base, *for the A-C market only*, are given for each carrier. Due to the addition of demand distribution, the chart is organized slightly differently. The labels for control method permutation run over the columns. Within each group of scenarios having a common RM assignment, the results for the three demand distributions appear in the order they were introduced. The derivative total market changes were suppressed because of space limitations, but appear below in tabular form.

The results for the leftmost FCFS/EMSRa scenarios, with LSC innovating, merit little comment. Under FCFS/EMSRa-1, LSC uses control to maximize revenue from the few passengers who cannot be accommodated on NSC, increasing revenues by 9%. This is achieved by selling up leisure requests, largely to M class, which has no effect on NSC revenues. For high local demand conditions, LSC's A-C revenues remain essentially unchanged. First, with the lower A-C demand, NSC does not refuse many requests, leaving LSC little demand to manage in the first place. Furthermore, the high local demand minimizes the seats available to prospective A-C passengers. Finally, a low local split

Demand/Control Scenario	% Carrier Revenue Change in A-C Over FCFS/FCFS			
	DF = 0.9		DF = 1.2	
	NSC	LSC	NSC	LSC
LSC Innovates				
Equal Split	+0.4%	+8.7%	-0.4%	+66.4%
High Local Split	+1.6%	-0.5%	+4.0%	-17.2%
Low Local Split	-0.7%	+17.7%	-0.7%	+106.7%
NSC Innovates				
Equal Split	+3.4%	+14.3%	+16.7%	+28.9%
High Local Split	-0.8%	-0.7%	+1.4%	+4.4%
Low Local Split	+16.7%	+11.8%	+57.6%	+57.4%
LSC & NSC Control				
Equal Split	+6.4%	+20.4%	+15.7%	+56.7%
High Local Split	+2.6%	-1.8%	+5.5%	-8.8%
Low Local Split	+17.6%	+20.0%	+76.4%	+75.8%

LSC = local service carrier
NSC = nonstop carrier

Table 5.6. Three-Market Scenario Revenue Effects by Carrier, A-C Market

raises the A-C revenue for LSC beyond that earned in an equal split. More, and higher-yield, passengers cannot find space on NSC and the lack of local traffic allows LSC to accommodate these requests profitably.

Figure 5.10 becomes interesting, and perverse at first glance, in the middle set of scenarios. For EMSRa/FCFS-e, the revenue benefit for the laggard LSC, nearly 15%, greatly exceeds that of the innovating NSC, just over 3%. Following NSC's innovation, LSC's A-C market revenue increases more in percentage terms than NSC's. Even when LSC alone had control (FCFS/EMSRa-e), it did not fare as well.

To understand why this occurs, Figure 5.11 has been created, summarizing the A-C traffic compositions for each carrier in the three equal split scenarios. Only Q class traffic is treated separately—the three higher-yield classes are aggregated to obtain a single measure of "quality" bookings. The first column in the right-hand half corresponds to FCFS/EMSRa-e. As indicated above, control affords an impressive fare class mix, with Q class constituting just half of the A-C bookings. Yet the total traffic is low since NSC, operating at FCFS, accepts every request not proscribed by advance purchase restrictions or capacity limitations (leftmost column of Figure 5.11). In EMSRa/FCFS-e, the next columns to the right, the situation is reversed. Over 90% of the LSC A-C bookings fall into Q class. But the critical point is that the rise in LSC A-C traffic volume, which occurs as NSC practices inventory control (second column, directly beneath NSC label), more than offsets this dilution. Dumping of longer haul A-C passengers can actually improve the individual market revenue performance of a local service carrier by displacing local market passengers. A localized "positive market externality" has been discovered, a rarity outside of microeconomic texts.

Returning to Figure 5.10, NSC innovation under a high local split again produces no revenue effects because of the low NSC load factor entailed by this market demand distribution. However, at the low local split, NSC matches the benefit LSC achieves when it is the innovator. Unlike FCFS/EMSRa-e, the rival also shares in the gains. At such a high effective demand factor for the NSC flight, traffic, albeit lower-yield traffic, continues to spill over to the less preferred A-C path.

In the last triplet of scenarios, with both carriers managing their inventory, the revenue changes at all three demand distributions can be closely approximated by adding together the results for the two analogous innovation scenarios. At an equal split, the spill over effect from NSC's control combine with the sell up potential from LSC's control to bring the advantage for LSC to 20%. Figure 5.11 demonstrates that both carriers do indeed take the best from each innovation world. LSC has higher total traffic than FCFS/EMSRa-e but without the diluted composition of EMSRa/FCFS-e. NSC nearly equals the total traffic count it observed as laggard, with the positive fare mix achieved as innovator. Meanwhile, a high local split continues to frustrate any attempt at RM revenue enhancement—the number of passengers in A-C is simply insufficient to make inventory control worthwhile. In EMSRa/EMSRa-l, NSC approaches the revenue gain of LSC by "skimming the cream" of the A-C passenger pool before passing the sizable remainder on to LSC.

While these findings are fascinating and quite unexpected, it should be stressed again that they have been developed from the A-C market results only. Returning to Table 5.5 and the earlier system-level discussion, the positive externality effects observed in A-C when LSC is the laggard do not translate to the system level. The extra A-C traffic can only be carried at the expense of local traffic and the two factors wash out, or perhaps net to be slightly negative.

* * *

The experiments in a stylized three-city network have provided a glimpse of the complexities which await the modeling of even a limited subnetwork in the current industry environment. By boiling down the multi-market problem to its barest essentials, it has been possible to trace the important traffic redistribution and associated revenue effects for various states of RM control. The system carrier revenue improvements echo the broad RM effects discovered in the first class of balanced single market scenarios. The basic game theoretic implications of inventory control—that carrier and market revenue benefits from RM depend on the set of control policies fielded by all competitors—have been well established. However, just as in the asymmetric single market scenarios, certain additional results, specific to the competitive setting, have been discovered for the three-market network. The highlights from the results for this class of competitive scenarios, as well as for the earlier single market tests, follow in the final chapter.

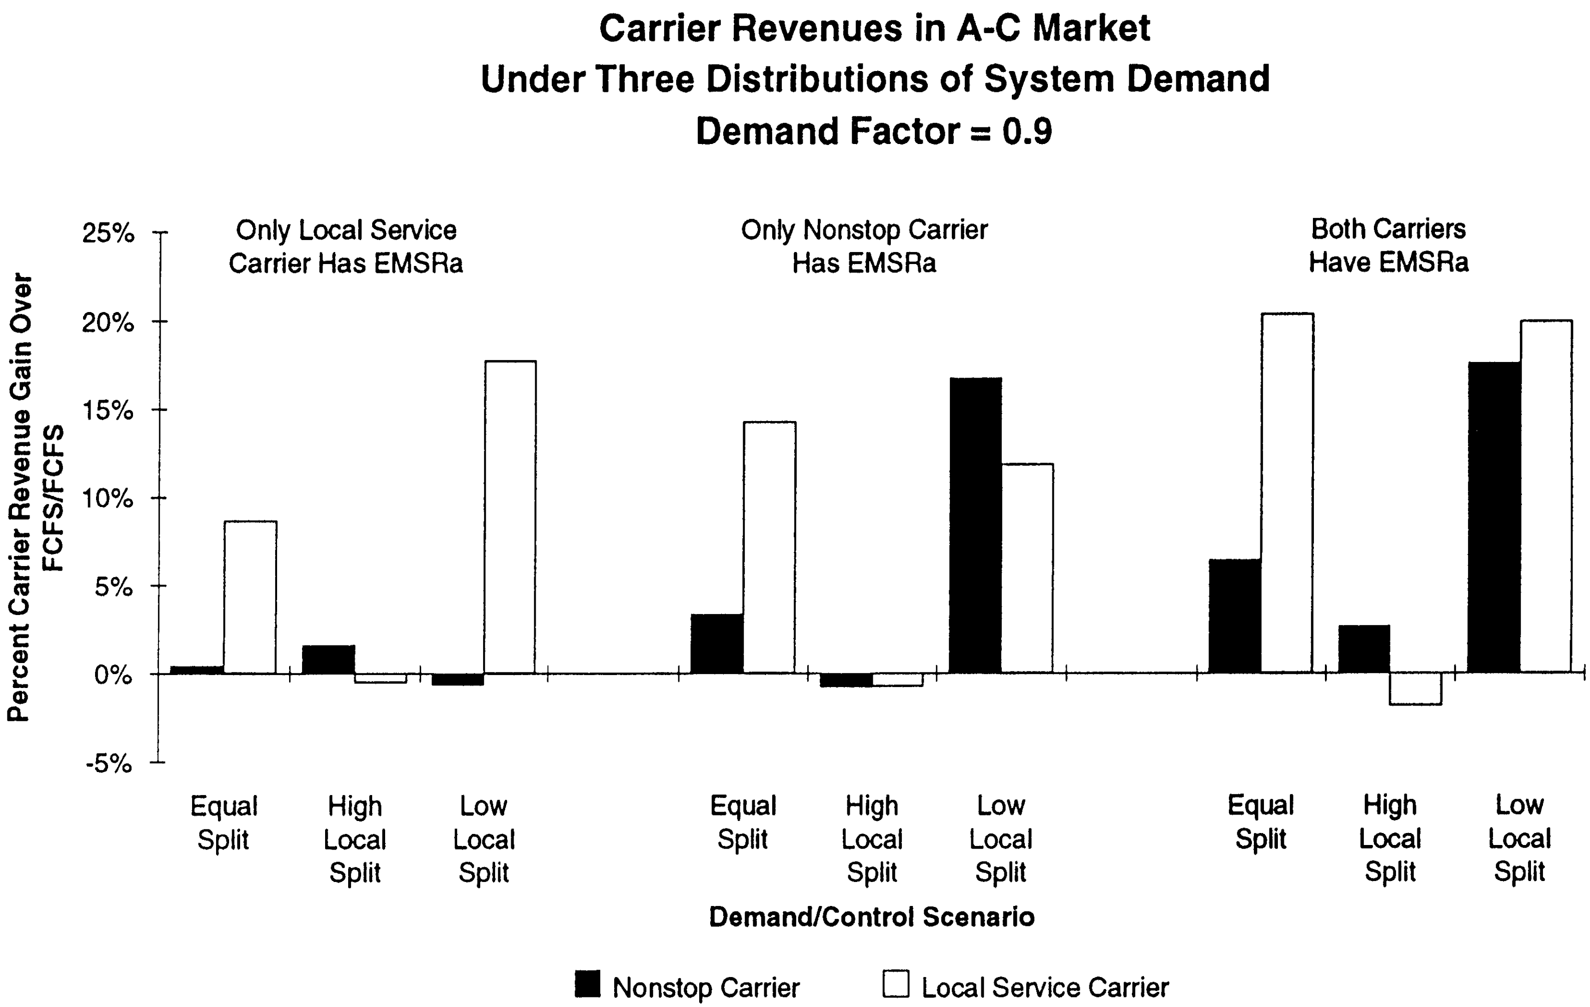

Figure 5.10. Value of Innovation, A-C Market, Demand Factor = 0.9

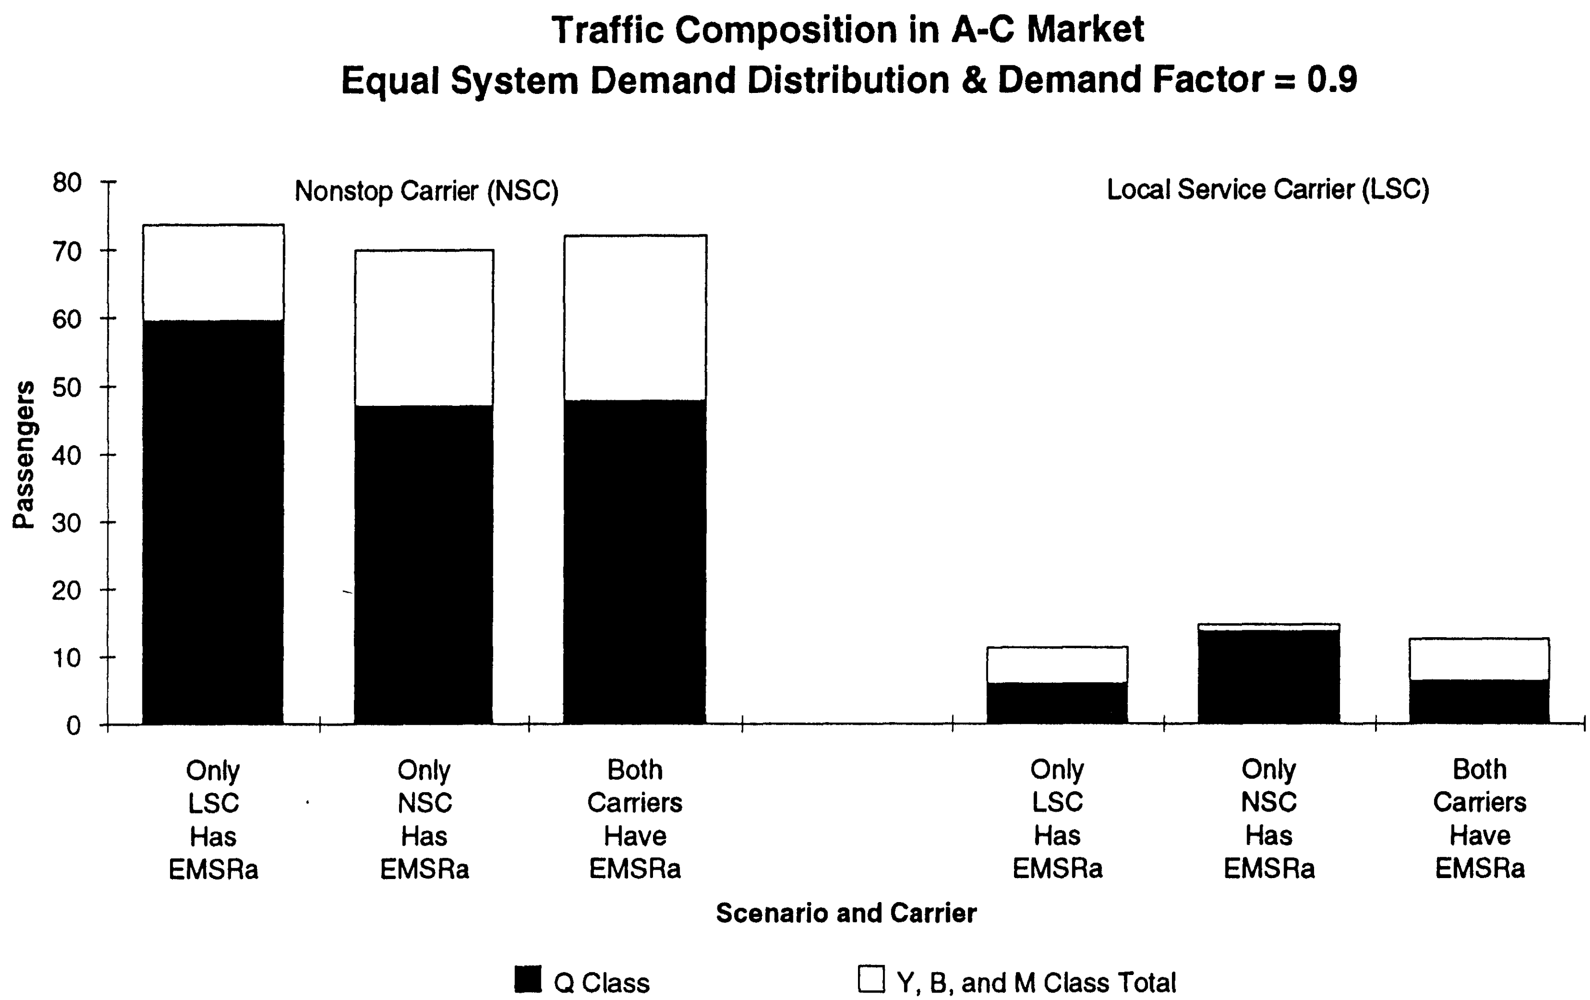

Figure 5.11. Traffic Composition, A-C Market, Demand Factor = 0.9

6 Conclusions

6.1 Summary of Simulation Results

The central premise of this thesis has been that the value of RM inventory control cannot be assessed independently of the competitive forces which determine airline performance. Previous booking simulations have effectively short-circuited the passenger choice process by assigning demand to an isolated flight and comparing the revenue benefits for alternate optimization algorithms given this captive demand profile. In contrast, the PODS / OCSE approach developed here recognizes that passengers in an OD market are free to choose among all available paths in the market. The competitive strength of each carrier serving the market derives from the level of service offered by the carrier, relative to its rivals. Conventional measures of competitive advantage have included frequency share and relative timing of departures. This thesis demonstrates that RM capability must also be considered a competitive advantage, interacting with all other competitive dimensions to produce a revenue benefit specific to the market setting.

To explore these interactions, three classes of competitive scenarios have been studied under various demand conditions. In the simplest form of competition, the symmetric two path scenarios, all competitive dimensions except RM are identical for the rival carriers. The second class, the asymmetric dominant carrier scenarios, tests the effects of frequency and schedule placement. Finally, expansion of the network to three cities introduces differing path qualities and the complication of multi-path legs. In each class of scenarios, simulation results have been presented to illustrate the revenue and traffic consequences from the adoption of RM. Both the total benefit to the industry and the distribution of this benefit among the individual carriers have been considered.

With respect to aggregate industry performance, the results for different RM combinations clearly indicate that the competition for market revenue is not a zero-sum game. Demonstrated in its purest form in the symmetric two path cases, the revenue gains of an innovating carrier do not come at the expense of the rival. Figure 5.1 shows that an EMSR innovator achieves a percentage improvement of 8-10% over the FCFS/FCFS base, while the rival suffers a penalty of less than 2%. Although the absolute sizes of the changes vary in each scenario, the total market always experiences a net revenue improvement when one carrier unilaterally adopts control.

The innovator's revenue dividend thus stems only partly from its ability to redirect discount requests to the uncontrolled alternative—reflected in the loss for the FCFS carrier. The remainder is provided by internal management of the fare class distribution. One of the assets of the modeling approach followed in this thesis is the relaxation of the common simulation assumption that passengers have an intrinsic fare class demand. Rather, given the specification of certain behavioral attributes, whose sensitivities have been documented in Chapter 4, passengers choose the best path / fare class combination among those still available when they appear in the booking process. Figures 5.2 and 5.3 capture the new fare class mix in the symmetric scenarios, made possible by variations in willingness-to-pay and attributed restriction cost attributes which, both by passenger type and from one individual to another within a passenger type.

Inventory control innovation thus does not obey the standard rules of competitive advantage. The benefit from inventory control derives mostly from encouraging passengers to buy up to higher-yield fare classes. Per its charter outlined in Chapter 2, RM enhances the price discrimination effort by selling capacity at fares nearer the *reservation price*[44] of each prospective passenger. Only a small fraction of the revenue improvement occurs from "stealing" the desirable traffic of a competing unmanaged carrier. Acquisition of control is a dominant strategy in that, regardless of the competitive situation, it is better to practice some form of control than remain at FCFS. Furthermore, an additional contribution of this thesis has been to demonstrate that the absolute improvement from a single optimization algorithm is not constant (as implied by previous simulation efforts), but rather decided both by the competitive setting and the collective control policies of rival carriers.

As a corollary of the non-zero-sum game result, the collective industry is found to be better off in revenue terms when both carriers have an equivalent control capability than when both operate under an FCFS discipline. In fact, in the two path symmetric scenarios, the recovery of the laggard carrier generates a revenue improvement equal to the innovation effect.

---

[44]A term from microeconomic theory representing the maximum price acceptable to an individual consumer.

This finding directs the discussion to the question of how the benefits from RM shift from one carrier to the other in the transitions between control scenarios. Over the range of competitive settings, three possibilities emerge for this transfer. As already noted, the innovator may observe no change in revenue when the rival acquires control. Alternatively, the innovator's first-mover advantage may be sharply curtailed by the laggard's recovery. This possibility is exemplified in the dominant scenarios by the evolution from FCFS/EMSRa-o to EMSRa/EMSRa-o (see Figure 5.6), as the revenue advantage achieved by the weak innovator of nearly 9% falls to just over 2% once the strong carrier implements a matching control policy. Finally, the three-city scenarios offer a third possibility. Although the system-level carrier revenue effects agree with the previous findings for the value of innovation, a special result has been discovered in an isolated analysis of the A-C connecting market. Specifically, NSC innovation can lead to a percentage improvement for LSC's A-C market revenue which is greater than the benefit realized by NSC itself. Figure 5.10 charts this result for the EMSRa/FCFS-e scenario. While this is a localized market phenomenon which does not carry over to system revenues, it is illustrative of the complex, and sometimes counterintuitive, results created by the interactions of RM and other competitive forces.

Analysis of the underlying traffic movements which produce these revenue effects also stresses the interdependence of the RM controls fielded by each carrier. Specifically, the fate of passengers spilled from a controlled flight is decided by the RM method employed by the rival carrier. Nowhere is this more apparent than in the three-city EMSRa/FCFS-e scenario, whose special revenue properties have been identified above. Figure 5.11 resolves the revenue result—when only NSC has EMSR, enough leisure passengers are spilled to LSC that the dilution of LSC's A-C market traffic is more than offset by the additional volume. This spill-over effect simply could not be represented in the simulation of an isolated flight—a spilled passenger would not have the chance to be reaccommodated.

Finally, Figures 5.4-5.6 concisely illustrate the revenue impact of each variant of EMSR control against a static RM capability, over a range of DFs—the single most important demand attribute to the success of RM. Beginning from a DF of 0.8, both EMSRa and EMSRb post gains of 3-5% regardless of the competitor's control method. At the other end of the DF spectrum, EMSRa and EMSRb are again indistinguishable, but the maximum revenue advantage now depends on the rival control policy—the 70% improvement over the FCFS/FCFS base attainable against FCFS falls to 50% against either EMSR. The results from the middle of the DF range almost single-handedly justify the efforts of this thesis. The difference between the revenue benefits of EMSRa and EMSRb is at its largest, about 10%, against a rival with FCFS at a DF of 1.0 (Figure 5.4). In contrast, against EMSRb (Figure 5.6), the two revenue trends over the entire DF range nearly parallel each other, with EMSRb edging EMSRa by just 3-4%. These results conclusively show that the benefits of different optimization algorithms vary when implemented against alternate rival control methods and under different DFs. Further original findings will undoubtedly emerge once additional scenario conditions (to be plotted along the x-axis) and control methods have been investigated in the competitive proving grounds established by PODS / OCSE.

6.2 Planned Research Extensions

To this end, MIT's Flight Transportation Laboratory and the Boeing Commercial Airplane Group intend first to expand the arsenal of RM options available in the modeling of competitive settings. As noted in Section 4.4, because of its modular architecture, PODS can theoretically simulate a formulation of any class of control method—from more sophisticated leg-based approaches to elementary virtual nesting and the newest network-control algorithms. Bearing in mind that PODS has been designed for a microcomputer platform, experiments with the more computationally intensive control methods would likely require execution under a scaled-down version of OCSE. Such a model reduction might be accomplished by eliminating a fare class or reducing the number of booking periods.

Similarly, the field of forecasting systems will be enlarged. Although a regression routine has already been coded, this option could not be evaluated in the thesis due to recent modifications of its logical structure. Once initial results have been validated, regression will join pick-up as a candidate forecasting method. A third forecasting approach, originally created for internal use at Boeing, may also be tested.

A third area for further research concerns the existing model of the reservation process. The simplifying assumptions listed in Section 3.1—namely the omission of cancellation, overbooking, and no-show effects—limit the applicability of current PODS results to real-world airlines. Incorporation of these harsh realities would be expected to lower the absolute revenue benefits attainable with basic RM innovation. By supposing a "perfect" booking process, the data in this thesis constitute an upper bound on the value of implementing RM. However, it remains to be seen whether the same relative results and trends hold under a more accurate representation of the reservation system. Airlines constantly wrestle with these confounding factors and there is no single solution which, for instance, neutralizes the no-show phenomenon. But even the roughest characterization of these effects would significantly further the current understanding of how revenue improvements from RM vary in uncertain demand / booking conditions.

Even within the scope of the existing RM alternatives and reservations model, many new experiments can be performed under more complicated and realistic competitive scenarios. These would explore the RM benefits in unbalanced single- and multiple-market scenarios which incorporate carrier-specific fares within a market, and a range of flight leg capacities. Additional carriers (perhaps with individual image factors) and cities may be modeled to flesh out a real-world subnetwork. Networks containing up to ten cities may reasonably be simulated with the current version of PODS.

To conclude, in the longer term, when a larger assortment of RM and forecasting methods has been compiled and the first revisions of the reservation model have been implemented, PODS / OCSE will become an even more powerful research tool. Following the scenario framework introduced here, the competitive landscape of actual market settings may be simulated in considerable detail. Existing and hypothetical inventory control approaches can then be pitted against each other, and the results compared in laboratory conditions—that is, allowing sensitivity analyses of all relevant behavioral and demand assumptions. While no model can replicate the empirical results yielded only by true revenue service, PODS achieves a quantum conceptual advance over earlier simulations by explicitly recognizing the implications of competition for the value of revenue management innovation. The experimental evidence from OCSE presented here eminently supports this hypothesis.

References

Alamdari, F. E. and I. G. Black. 1992. Passengers' Choice of Airline Under Competition: The Use of the Logit Model. *Transport Reviews.* **12**, 153-170.

Belobaba, P. P. 1987. Air Travel Demand and Seat Inventory Management. MIT Flight Transportation Laboratory Report R87-7.

Belobaba, P. P. 1992a. Optimal vs. Heuristic Methods for Nested Seat Allocation. Presented to AGIFORS Yield Management Study Group, Brussels, Belgium.

Belobaba, P. P. 1992b. The Revenue Enhancement Potential of Airline Yield Management Systems. *ASTAIR Proceedings*, 143-164.

Belobaba, P. P. 1994. Network Seat Inventory Control Without Network Optimization. Presented to AGIFORS Yield Management Study Group, Hong Kong.

Ben-Akiva, M. and S. R. Lerman. 1985. *Discrete Choice Analysis: Theory and Application to Travel Demand.* MIT Press, Cambridge, MA.

Berge, M. E. and C. A. Hopperstad. 1993. Demand Driven Dispatch: A Method for Dynamic Aircraft Capacity Assignment, Models and Algorithms. *Operations Research.* **41**, 153-168.

Boeing Company. 1988. *Domestic Fare Survey.*

Boeing Commercial Airplane Group. 1993. Decision Window Path Preference Methodology: Time Mode. AMADWM-1.

Curry, R. E. 1990. Optimum Airline Seat Allocation with Fare Classes Nested by Origins and Destinations. *Transportation Science.* **24**, 193-204.

Hansen, M. 1990. Airline Competition in a Hub-Dominated Environment: An Application of Noncooperative Game Theory. *Transportation Research*, **24B**, 27-43.

Hopperstad, C. 1995. Potential Application of the Boeing Passenger Origin / Destination Simulator [PODS] to Yield Management Simulation Planning. Boeing Commercial Airplane Group.

Smith, B. C., *et al.* 1992. Yield Management at American Airlines. *Interfaces.* **22**, 8-31.

Tan, J. C. 1994. Leg-Based Heuristic Methods to Network Seat Inventory Control. MIT Flight Transportation Laboratory Report R94-7.

Waldman, G. L. 1993. A Study of the Practicality and Profit Enhancement Potential of Demand Driven Dispatch in Airline Hub Operations. MIT Flight Transportation Laboratory Report R93-3.

Weatherford, L. R. and S. E. Bodily. 1992. A Taxonomy and Research Overview of Perishable-Asset Revenue Management: Yield Management, Overbooking, and Pricing. *Operations Research.* **40**, 831-844.

Williamson, E. L. 1992. Airline Network Seat Inventory Control: Methodologies and Revenue Impacts. MIT Flight Transportation Laboratory Report R92-3.