Modeling Multiple Airport Systems:
A Positive Feedback Approach

Mark Hansen
Qiang Du

Working Paper
UCTC No 404

The University of California
Transportation Center

University of California
Berkeley, CA 94720

The University of California Transportation Center

The University of California Transportation Center (UCTC) is one of ten regional units mandated by Congress and established in Fall 1988 to support research, education, and training in surface transportation The UC Center serves federal Region IX and is supported by matching grants from the U S Department of Transportation, the California Department of Transportation (Caltrans), and the University

Based on the Berkeley Campus, UCTC draws upon existing capabilities and resources of the Institutes of Transportation Studies at Berkeley, Davis, Irvine, and Los Angeles, the Institute of Urban and Regional Development at Berkeley and several academic departments at the Berkeley, Davis, Irvine, and Los Angeles campuses Faculty and students on other University of California campuses may participate in Center activities Researchers at other universities within the region also have opportunities to collaborate with UC faculty on selected studies

UCTC's educational and research programs are focused on strategic planning for improving metropolitan accessibility, with emphasis on the special conditions in Region IX Particular attention is directed to strategies for using transportation as an instrument of economic development, while also accommodating to the region's persistent expansion and while maintaining and enhancing the quality of life there

The Center distributes reports on its research in working papers, monographs, and in reprints of published articles It also publishes *Access*, a magazine presenting summaries of selected studies For a list of publications in print, write to the address below

University of California Transportation Center

108 Naval Architecture Building
Berkeley, California 94720
Tel 510/643-7378
FAX 510/643-5456

Modeling Multiple Airport Systems: A Positive Feedback Approach

Mark Hansen
Qiang Du

Institute of Transportation Studies
University of California
Berkeley, CA 94720-1720

Working Paper
October 1993

UCTC No 404

The University of California Transportation Center
University of California at Berkeley

PREFACE AND ACKNOWLEDGEMENTS

This report documents research performed for the State of California, Department of Transportation, under contract #65H998 - MOU 79.

The contents of this report reflect the views of the authors, who are responsible for the facts and accuracy of the data presented herein. The contents do not necessarily reflect the official views or policies of the California Department of Transportation. This report does not constitute a standard, specification, or regulation.

EXECUTIVE SUMMARY

This report proposes and applies a model of traffic allocation in multiple airport systems (MAS). Unlike most previous studies, which focus on a passenger's choice of airport for a given set of service attributes (frequency, fare, etc.), this model assumes that service attributes are endogenous to the system, and directly related to airport traffic volume. Thus a positive feedback exists whereby an airport becomes more attractive, the more traffic it has. The distribution of traffic in a MAS reflects the combined effect of travelers choosing airports closer to their trip origin and the positive feedback effect. Airports located more conveniently to the market have an initial advantage stemming from their greater accessibility, which is then amplified by positive feedback.

The model is applied to the San Francisco Bay Area, which is served by three major commercial airports--Oakland, San Francisco International, and San Jose. We first estimate an airport choice model in which airport destination market traffic, total airport traffic, and travel time to the airport are included as attributes. The model yields statistically significant coefficients of expected signs. We then use the calibrated logit model, along with data concerning the origin of air passenger trips in the Bay Area, by destination market, and ground access travel times to each airport to predict equilibrium market shares. We find good overall agreement between the predicted shares and observed shares for markets. The model predictions are most accurate for markets with large traffic volumes, and more accurate for Oakland and San Francisco airports than for San Jose. Statistical tests reveal that there is a significant relationship between predicted and actual values, and that if equilibrium share is treated as an explanatory variable, it accounts for most of the variation in actual share in large markets. Finally, predicted market share is found to be a statistically significant predictor of observed seat share, supporting the assumption that airport service attributes are to a large extent endogenous.

The calibrated model is used to predict the market share of Buchanan field, an airport in Contra Costa county where commercial service has been offered sporadically. Predicted shares of traffic to Southern California are in the range of 10 to 25 per cent, although actual market share never exceeded 3 per cent. We consider several possible explanations for this disparity. One is that Buchanan's directors did not support expanding and improving the airport so that it could better accommodate commercial activity. In addition, there is evidence that many air travelers were unaware of the services available out of Buchanan. These explanations suggest that our model is applicable only to established airports that are committed to providing commercial service.

Our main conclusion is that the positive feedback process upon which our model is based plays a large role in determining the distributions of commercial traffic in a multiple airport region. Furthermore, the key exogenous factor in this process, the regional distribution of air trip origins, is found to vary significantly from one destination market to another. Thus good traffic forecasts require market-specific predictions of these distributions. There is a pressing research need to develop models that allow such predictions. In addition, analysis of other multiple airport systems is required to assess the transferability of the model, and also to improve our understanding of how other factors not directly in the model, such as capacity limitations, hubbing, and accessibility by non-automobile modes, affect traffic allocation. With such improvements, planners and policy makers will be able to use the model to determine how to most effectively balance the competing goals of accessibility, service quality, and infrastructure cost in multi-airport systems.

TABLE OF CONTENTS

LIST OF FIGURES

Page

LIST OF FIGURES (cont.)

LIST OF TABLES

Page

Modeling Multiple Airport Systems: A Positive Feedback Approach

1. Introduction

Many of the world's largest cities are served by more than one commercial airport. Such multiple airport systems (MAS) have recently been the object of increased attention, for several reasons. From an academic viewpoint, multiple airport systems are interesting because they provide a window on the preferences of air travelers -- particularly the relative values they attach to ease of access to the airport versus the quality (and cost) of service from it -- and the competitive strategies of airlines. From a practical standpoint, understanding multi-airport systems is important so that the consequences of various transportation supply actions -- surface transportation improvements, traffic caps or service restrictions on "over-utilized" airports, creation of new airports, or investments in existing ones -- can be assessed.

By far the most widely studied aspect of MAS is the air traveler's choice of airport. Studies by Harvey (1987), Ashford (1989), and others have used data on passenger choice of airport to explain how flight frequency, service directness, fare, and accessibility influence these decisions. These models have generally been successful, at least in the sense of yielding plausible, statistically significant, coefficients that capture the relative weights travelers place on different service attributes. But while these studies contribute to the understanding of travel behavior, they cannot by themselves explain or predict the behavior of an MAS, since they treat the supply-side of the system as exogenous.

To date, efforts to incorporate both the demand and supply sides of multi-airport systems in a single framework have been limited. The Multiple Airport Demand Allocation Model, or MADAM, is the best known effort to produce an operational model in this area. Other models, such as that offered by de Neufville and Gelerman (1973), are intended to highlight general tendencies of MAS. We will argue below that neither of these models is adequate for understanding or predicting the behavior of MAS.

This report proposes a new model of MAS. The model treats an MAS as a <u>positive feedback system</u> in which an airport's attractiveness to passengers is positively related to the volume of its passenger traffic. The model presented is simple, easy to work with, and based on plausible behavioral assumptions. It requires relatively little input data, and yields accurate predictions of airport market shares in the San Francisco Bay Area, the MAS to which the model is applied.

In addition to its methodological contribution, this report shows that, in the Bay Area at least, the distribution of air trip origin locations varies across destinations, and that this variation explains a large amount of the observed destination-to-destination variation in airport market shares. The destination-to-destination variation in the relative accessibility of the Bay Area airports is modest in itself, but when amplified by the positive feedback effect, it has a pronounced effect on airport market shares.

The remainder of this report is organized as follows. Section 2 contains a brief literature review. Section 3 presents the basic structure and concepts of the proposed positive feedback model. Section 4 describes model development and assesses model performance. Section 5 uses model results to analyze service supply, while in Section 6 the model is used to predict market share for a fourth Bay Area airport. Section 7 offers conclusions.

2. Literature Review

Prior models of multiple airport systems fall into two main categories. The first category focusses on the choice of airports by air travelers. Studies of this type examine how air travelers choose between a set of airport alternatives, based on (1) the services provided at the airports and (2) ease of ground access to the airports. These factors are treated as exogenous. Thus, to use such models predictively, one must make assumptions about airport service and ground access characteristics. The second category of multiple airport system models treats certain airport service characteristics, most notably flight schedules, as endogenous to the system. Obviously, this reduces the assumptions required for predictive use of the models.

The first airport choice studies appeared in the mid-1970s. Skinner (1976) developed a disaggregate airport choice model for the Baltimore-Washington region. Multinomial logit models including an airport level-of-service term, based on flight frequency, and a ground access level-of-service term were estimated. Different measures of both airport and ground access level-of-service were considered, and separate models for business and non-business trips were developed. Both level-of-service terms were found highly significant, and acceptable goodness-of-fit measures were obtained. Model performance did not vary significantly with respect to the specific level-of-service measures used.

Kanafani et al. (1977) developed an aggregate choice model of airport choice, based on data for the California corridor market between San Francisco and Los Angeles. Since both of these regions have more than one airport, analysis was based on the choice of airport on both ends of the trip. Observed market shares of traffic between different subareas of the two regions were related to total estimated travel time and flight frequency for the various airport pairs. The estimated coefficients were statistically significant and of expected sign, and the calibrated model fit observed data well.

More recently, Harvey (1987) estimated an multinomial airport choice model on data for the San Francisco Bay Area. The main differences between the Harvey and Skinner models are that the former includes airport dummy variables to capture airport-specific factors that influence attractiveness to travelers, and uses non-linear transformations of the travel time and frequency variables. In the case of travel time, Harvey found decreasing marginal disutility, while for flight frequency he found decreasing marginal utility.

Innes and Doucet (1990) studied airport choice of travelers originating from the northern half of New Brunswick, Canada. They employed a binary logit model. Unlike the earlier studies, they did not consider flight frequency, but rather aircraft type, service directness, and flying time as airport level-of-service variables. For ground access level of service they tried a number of variables, all defined on the basis of distance from trip origin to airport. Their estimation results suggested, implausibly, that distance had a positive effect on utility, probably because of the way in which they constructed the airport pairs used in the analysis. The airport level-of-service variables were of the expected sign, however, with aircraft type (jet versus non-jet) the most important.

Ashford and Benchemam (1987) developed a choice model for Central England. Unlike the other studies, this one included fare along with flight frequency as an airport level-of-service variable. The fare coefficient was negative and statistically significant for international leisure and domestic travelers, but positive for the international business and inclusive tour market segments. Frequency and travel time were statistically significant and of expected sign for all travel classes.

Table 1 presents a summary of comparable results from the Skinner, Kanafani, Harvey, and Ashford models. The Harvey and Ashford studies yield consistent travel time coefficients that are about twice the magnitude of those obtained by Skinner. The studies are less consistent with respect to the frequency coefficients. In addition to underlying differences in traveler preferences, these

Table 1.
Coefficient Estimates, Airport Choice Models

Variable	Skinner (1976)	Kanafani (1977)	Harvey (1987)	Ashford (1987)
Travel Time (min)	-0 06 to -0 08	-0 10	-0 12 to -0 17	-0.14 to -0 23
Frequency (per day)	0 09 to 0 12	0.003 to 0 002	0 to 0 5[1]	1.07 to 2.70
ρ^2	0 59 to 0 69	--	0 61 to 0 77	0.84 to 0 92

Note
1 Estimated from Figure 4 of Harvey (1987) The lower value applies when frequencies exceed 8 per day, and the higher values where daily frequencies are less than 6

disparities probably reflect the inappropriateness of including a linear frequency term in the utility. Such a term implies that an additional flight has the same impact on utility, whether it increases frequency from 1 to 2 per day or 10 to 11 per day. Both Skinner and Harvey estimate models that incorporate a non-linear relationship between flight frequency and utility, but Ashford does not. To maintain comparability, Table 1 is based on the versions of the Skinner model with a linear frequency term. Since Harvey does not consider such a model, we present derivatives of utility with respect to frequency for "high" and "low" values of this variable as the linear coefficients.

As already noted, understanding travelers' airport choices is necessary but not sufficient to understand the behavior of multiple airport systems. A few studies have attempted to go further by considering both supply-side (airline) and demand side (traveler) behavior in such systems. de Neufville and Gelerman (1973) were among the first to attempt this. They argued that commercial air traffic will inevitably concentrate at one dominant airport in a multiple airport system. The basic reason for this is the "s-curve effect," by which an airline or airport with a high (low) share of flights in a given origin-destination market will attain an even higher (lower) traffic share. Using a simple game theory model, de Neufville and Gelerman argue that these effects result in competitive equilibria in which each airline concentrates its flights at one airport. Further, if one airport has a locational advantage, they claim that each carrier will select this airport as the one at which to concentrate. Thus, in any multiple airport system, a single airport will tend to dominate, with only a few flights (resulting from "second order effects," according to the authors) being offered from other airports.

de Neufville and Gelerman suggest a model of airport (airline) choice that captures the s-curve effect. It is (p. 542):

$$MS_A = \frac{FS_A{}^K}{\sum_i FS_i{}^K}$$

Where:

MS_A is airport A's market share;

FS_A is its frequency share;

K is some exponent (which must be greater than 1 in order to explain the s-curve effect).

Unlike the choice models documented above, this one considers the region as a whole rather than a particular origin point in that region. However, an origin-specific model embodying this postulated relationship can be specified by using the logarithm of frequency in the logit utility function, as Hansen (1990) does in his route choice model. Unfortunately, none of the origin-specific airport choice models reported in the literature employ this form, making it difficult to determine whether these models imply an s-curve effect similar to that postulated by de Neufville and Gellerman.

Even if the s-curve relationship is correct, it is apparent that the de Neufville/Gellerman theory is incomplete, because it gives so little weight to airport accessibility. This is best seen from the fact that the argument made to support the theory are not in any way tied to the size of the region. If their argument implies an urban area will tend to have one dominant airport, so also does it suggest that a single airport will dominate a state, country, continent, or larger geographical unit. Clearly, there is a point at which proximity to an airport outweighs the availability of more frequent service at other airports. Gellerman and de Neufville offer no convincing reason why, if more than one airport can exist in the United States, this cannot also occur in an urban area.

In light of the above, it is evident that models of MAS must be based on the specific geography of the region, particularly the location of trip origins relative to the airports. The most well known model that does this is the Multiple Airport Demand Allocation Model, or MADAM (National Capital Region Transportation Planning Board (NCRTPB), 1985). The model works as follows:

1. Local passenger demand is allocated to airports, based on the shortest door-to-door round trip ground access time from local zone of origination. Passenger demand is adjusted upward to account for non-local traffic.

2. Seats to be provided in each city-pair market from each airport are calculated based on the airport's passengers in that market and typical load factors. A scheduling table determines the number of flights to be offered, by size class, as a function of the seat volume and stage length of the city pair market. A heuristic is used to schedule the flights, based on a distribution of preferred departure times.

3. An average waiting time for each airport and city-pair market is calculated. This time, multiplied by a weighting factor, is added to the door-to-door round trip time. An airport-specific "facilitation time" may also be added.

4. Passenger demand is reallocated, as in Step 1, except that total time calculated in Step 3, rather than the ground access time, is used.

5. The process continues until an equilibrium, or the specified number of iterations, is reached.

6. If, after Step 5, there are airports whose passenger capacity is exceeded, passengers are removed and reallocated to the other airports that are under capacity.

The MADAM model is "calibrated" using an airport-specific ground access time correction and the waiting time multiplier (see Step 3, above). These variables are set so that the model predictions "replicate base year conditions as closely as possible (NCRTPB, 1985, p. 44)." When applied to the Baltimore-Washington area for 1981/82, the calibrated model predicted airport traffic shares, both by local origin and external destination, that matched observed data fairly closely.

There are numerous flaws with the MADAM model. The model has weak behavioral underpinnings. The "all-or-nothing" method for allocating passenger traffic from a given local origin to a given external destination is implausible. The scheduling table used to simulate the supply-side of the system is, at best, a rough approximation of empirically observed patterns of airline behavior, without any theoretical basis whatsoever. The model allocates traffic based on its non-stop, rather than its final, destination, and is therefore unable to deal with competition between non-stop, direct, and connecting services. Indeed, non-stop traffic levels depend on the outcome of such competition, since this will determine how much of the traffic goes non-stop to final destinations versus intermediate airports. Thus the demand variables that MADAM assumes to be exogenous are actually endogenous to a large extent.

In sum, the bulk of research in MAS has concerned the air traveler's choice of airport. While these studies have for the most part yielded plausible and statistically robust results, they do not address the supply-side of the system and are thus of little predictive value. The few efforts to address both demand and supply have not been convincing, whether because of their failure to adequately account for the role of airport accessibility, the implausibility of their behavioral assumptions, or failure to correctly distinguish exogenous and endogenous variables.

3. Proposed Model

The proposed model is based on three propositions concerning the preferences of air travelers in a MAS. First, travelers prefer airports that are closer to their trip origins. Second, travelers prefer airports with higher levels of traffic in their market with the strength of this effect increasing with market distance. (Throughout this paper, we use the term "market" to mean true origin-destination market, defined by the communities where an air trip itinerary begins and ends). Third, travelers prefer airports with higher levels of total traffic. The first of these propositions is obvious and self-explanatory. The second and third propositions require more explanation.

The second proposition is based on the existence of economies of scale at the origin-destination market level in air transportation. As traffic levels to a particular destination increase, costs per passenger associated with both airline and air traveler-supplied inputs decrease. These economies have several sources. First, load factor increases distribute the fixed costs of aircraft operation over a large traffic base. Second, frequency increases improve schedule convenience. Third, aircraft size increases may result in lower unit operating costs for airlines and a more comfortable cabin for passengers.

As a result of these economies, when traffic is too low, scheduled non-stop service is not economically viable. As traffic levels increase, therefore, so does the probability that non-stop service will be available. Since passengers strongly prefer such service (Hansen, 1990), a strong feedback effect is expected when traffic increases result in non-stop service becoming available. Once non-stop service is available, further increases in traffic increase attractiveness via the frequency and aircraft size effects. In addition, there may be a fare reduction effect, resulting either from reductions in airline costs or increases in competition as more carriers can "fit" into the market.

We expect market distance to strengthen the positive feedback effect, for two reasons. First, smaller aircraft tend to have shorter ranges. Thus, the minimum traffic necessary for non-stop service to be economically viable increases with distance. Second, service frequencies in long-haul markets tend to be lower, both due to larger aircraft and higher load factors (Baily, Graham, and Kaplan, 1985). Since there are diminishing returns to frequency in its effect on service attractiveness, there should be a greater gain in service attractiveness from a given percentage gain in traffic in a long-haul market.

When non-stop service is unavailable or unattractive, travelers may use connecting or multi-stop service. Even in this case, market traffic may have some impact on service quality. For example, shorter layover times may be available when the number of connecting passengers in a market is larger. However, we also expect traffic in other markets to affect the quality of this type of service, as indicated by the third proposition above. At the airport level, more traffic translates either into more destinations to which service is available, or more average traffic to each destination. The first effect will increase the number of connecting routings available, while the second -- by virtue of the scale effects discussed earlier -- makes the services on existing routings more attractive. Conceivably, increased airport traffic may also have other effects, both positive -- improved ground access, better airport services -- and negative -- greater congestion and walking distances. On balance, however, we expect the effect to be positive.

To operationalize the above propositions, we propose a logit model that predicts the probability that a traveler whose trip to destination k originates from location i within the multiple airport region will select airport j. The proposed model is:

$$P(j|i,k) = \frac{e^{V_{ijk}}}{\sum_{i} e^{V_{ilk}}} \tag{1}$$

$$\begin{aligned} V_{ijk} = {} & \alpha \cdot \log(PAX_{jk} + \theta \cdot (NLPAX_j + \sum_{n \neq k} PAX_{jn})) + \\ & \beta \cdot \log(DIST_k) \cdot \log(1 + PAX_{jk}) + \psi \cdot ATIME_{ij} + \gamma_j \end{aligned} \tag{2}$$

where:

PAX_{jk} is total origin and destination passengers from airport j to destination k;

$NLPAX_j$ is total nonlocal passenger traffic enplaning or stopping over at airport j;

$DIST_k$ is the distance of destination k from the MAS;

$ATIME_{ij}$ is ground access travel time (or generalized cost) between location k and airport i;

γ_j is the utility constant for airport j

The market and airport traffic terms are included in the argument of a single log function. The log transformation is used for two related reasons. First, we expect diminishing returns -- as traffic increases, its marginal impact of utility decreases. Second, insofar as the traffic terms reflect the "size" of each alternative -- the number of flights, airlines, and routings available -- the theory of alternative aggregation in logit models indicates that they should be incorporated in a logarithmic form. Both traffic variables are included within a single log term because this accords with our intuition about how the choice probability should behave as the market traffic term approaches zero. Specifically, as long as there is some traffic (and therefore some service) going to other destinations, then there will be connecting possibilities resulting in some finite probability that a passenger will select a particular airport to go to a given destination, even if there is no other traffic going to that destination.

Airport traffic includes both non-local and local traffic. For present purposes, we treat the non-local traffic as an exogenous variable, although in reality there is a simultaneous relationship between local and non-local traffic levels. To make non-local traffic endogenous, we would have to model a much larger system, including airports outside the MAS that compete with it for non-local traffic. Furthermore, non-local traffic depends on airline hubbing decisions of a strategic nature that are very difficult to predict. Thus, treating non-local traffic as exogenous is a necessary compromise to keep our modeling task tractable.

The second term in the utility function reflects the role of market distance in strengthening the positive feedback effect. Distance is included in logarithm form because we expect its impact to fall off as distance increases. Traffic is included in logarithmic form for the same reasons as it was in the first term. Since our arguments concerning the effect of distance are based on market, as opposed to airport, traffic, only the former is included. However, we again want to allow the possibility of a passenger choosing an airport for travel to a given destination even when its traffic in that market is zero. Thus we add the value 1 to market traffic.

The third term in the utility function reflects the difficulty of travel from zone i to airport j. The fourth term is an airport-specific constant that captures the utility (or disutility) associated with airport attributes not otherwise accounted for. Such factors include the levels of landside, terminal, and airside congestion, walking distances, and parking costs. In addition, the constant captures any airport-to-airport differences in supply-side behavior that result in certain airports having more attractive service for a given traffic level.

Although the airport-specific constant may capture congestion effects, the model does not explicitly include capacity constraints or the tendency of airports to become less attractive as traffic levels approach capacity. These are clearly limitations that must be addressed in future research. For the moment, the applicability of our model is restricted to MAS where capacity limitations are not an overriding factor in determining traffic allocation. The San Francisco Bay Area, to which we apply the model in this report, meets this criterion.

For the MAS to be in equilibrium, the market passenger levels must satisfy the following equations:

$$PAX_{jk} = \sum_{i} PAX_{ik} \cdot P(j|i,k) \quad \forall j,k \tag{3}$$

PAX_{ik} is total origin and destination passengers to destination k originating from location i within the region.

The number of these equations is the product of the number of airports in the MAS and the number of destinations to which people travel from the MAS. There are an equal number of unknowns -- the PAX_{jk}.

To illustrate the model, we consider a two-airport system in a linear city of unit length along which air trips, all to the same destination, are generated at a uniform rate, as shown in Figure 1. The distance between the airports is also of unit length, but the airports are offset from the ends of the city by a distance δ. When $\delta>0$, Airport 2 has a locational advantage over Airport 1.

According to the proposed model, passengers originating from a given point in the city will choose between the two airports based on their relative traffic levels and ground access time. Assuming ground access time is proportional to access distance, the model (stated in terms of the Airport 1 market share) becomes:

$$P(1|x) = \frac{e^{\alpha MS_1 - \beta (x+\delta+\frac{1}{2})}}{e^{\alpha MS_1 - \beta (x+\delta+\frac{1}{2})} + e^{\alpha (1-MS_1) - \beta (x-\delta-\frac{1}{2})}} \tag{4}$$

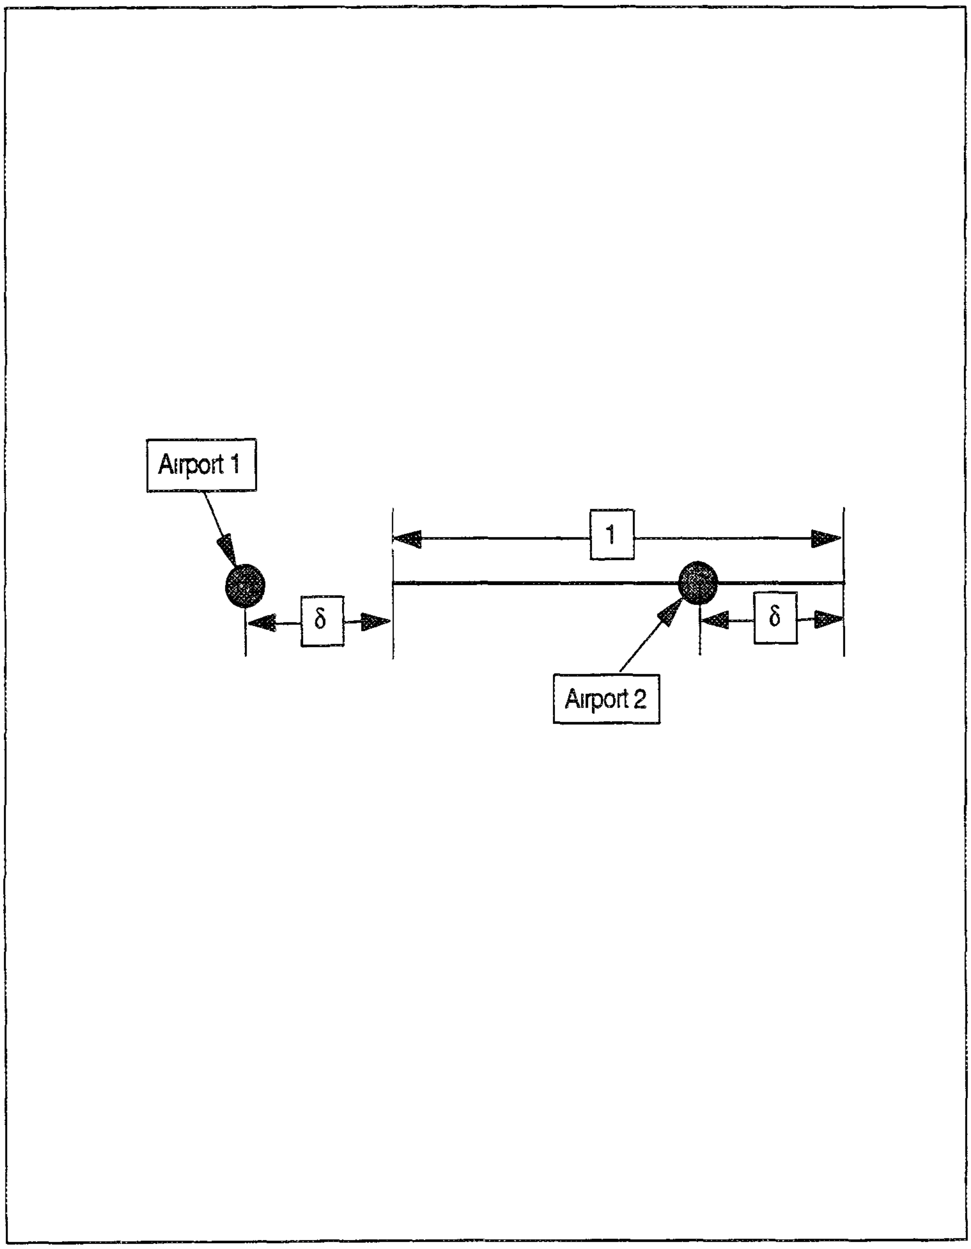

Figure 1. Geometry for Two-Airport System

$$MS_1 = \int_{x=-1/2}^{x=+1/2} P(1|x)\,dx \tag{5}$$

where:

MS_1 is the market share of airport 1;

x is the location of the passenger's origin, measured from the center of the city.

Equations 4 and 5 can be solved for the equilibrium market share for Airport 1.

Figure 2 depicts the equilibrium values of MS_1 for $\delta=0$ and three different sets of α and β values. The curves in Figure 2 depict the relationship between market share, treated as an independent variable, and the market share predicted by equations 4 and 5. The 45-degree line is the set of points where the two market shares are equal. Since both axes measure the same variable, only points on this line are feasible. Thus, equilibria occur when the line intersects with the curve. In this example, equilibria occur at $MS_1=0$, $MS_1=0.5$, and $MS_1=1$. The first and third of these will always be solutions, regardless of the values of α, β, δ. The second solution reflects the symmetry of the system when $\delta=0$. When $\delta\neq0$, this solution may either be shifted to another MS_1 value, or cease to exist.

Equilibria may be stable or unstable. If the curve cuts the 45-degree line from above as MS_1 increases at the intersection point, the equilibrium is stable. Such a situation exists for the $MS_1=0.5$ solution when $\alpha=0.5$, for either $\beta=1$ or $\beta=3$. The stability is illustrated for the case of $\beta=1$. Suppose the system is in disequilibrium, with $MS_1=0.7$. This leads to a new MS_1 of just over 0.6 -- closer to the 0.5 equilibrium value. On the other hand, consider the case where $\alpha=2$ and $\beta=1$. In the situation illustrated, a initial disequilibrium market share of 0.4 leads to a new market share of 0.26, further from the 0.5 equilibrium, indicating that $MS_1=0.5$ equilibrium is unstable. One can also see that when the 0.5 equilibrium in stable, the 0 and 1.0 equilibria are unstable, and

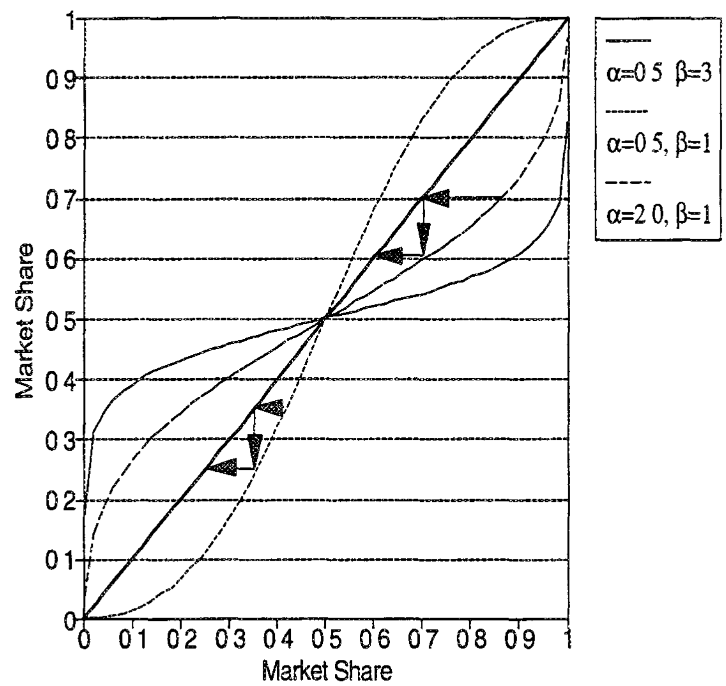

Figure 2. Equilibria for δ=0

vice versa.

When $\delta>0$, Airport 2 has a locational advantage, which the positive feedback effect tends to magnify. This is illustrated in Figure 3, where it is assumed that $\delta=0.1$ and $\beta=1$. When α is small, a stable equilibrium at $0<MS_1<0.5$ exists. As α increases, this equilibrium disappears, and for a certain range the only stable equilibrium is $MS_1=0$. As α continues to increase, however, an unstable equilibrium at $0.5<MS_1<1$ appears, and the previously unstable equilibrium at $MS_1=1$ becomes (locally) stable. Thus, if positive feedback is strong enough, an pre-existing airport in a less convenient location could maintain its position against a new, better situated, airport. If the feedback is slightly less strong, all traffic would migrate to the new airport, while an even weaker feedback effect will result in a division of the market between the two airports.

The effect of varying the locational advantage of Airport 2 is illustrated in Figure 4. As δ increases, the market share curve for airport 1 is pushed downward, resulting in a reduction of Airport 1's equilibrium market share. When δ increases beyond a certain point, the intermediate equilibrium points disappear, leaving only the stable equilibrium at $MS_1=0$ and the unstable one at $MS_1=1$.

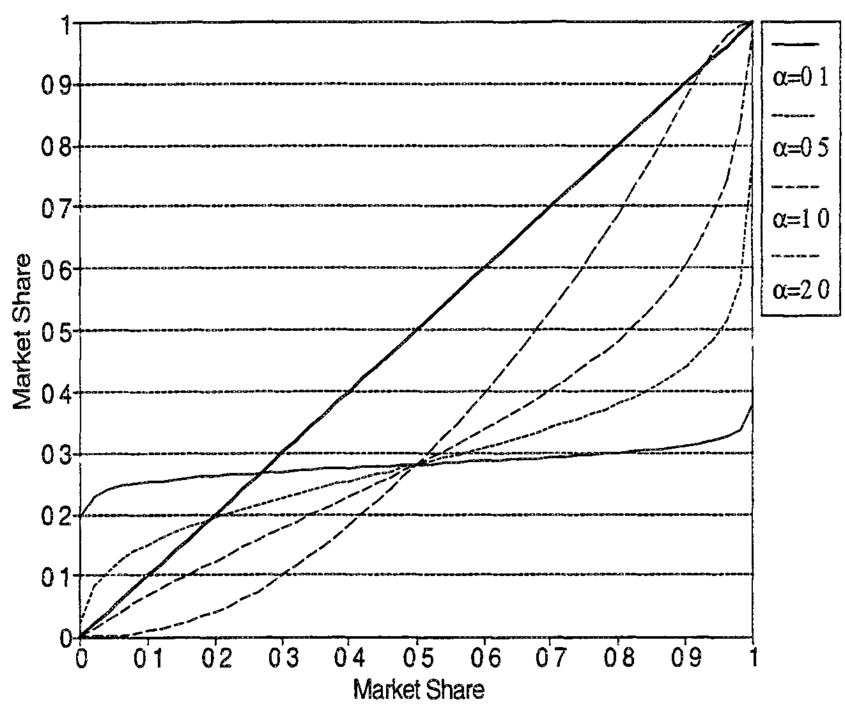

Figure 3. Equilibria for δ=0.1 and ß=1 under Alternative α Values

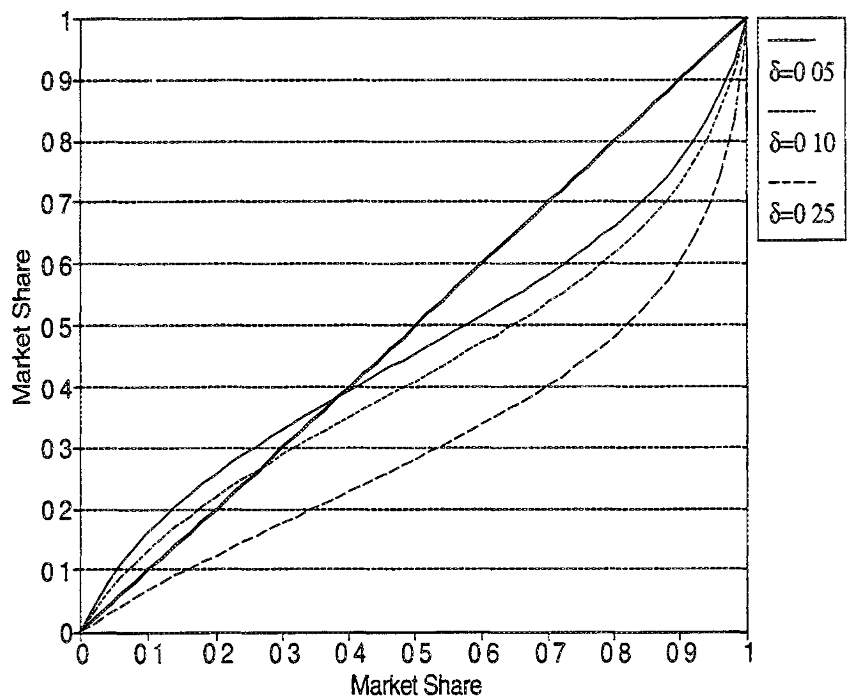

Figure 4. Equilibria for Alternative δ Values

4. Application to the San Francisco Bay Area

4.1 Background

The positive feedback model was applied to the San Francisco Bay Area. The Bay Area is served by three commercial airports -- San Francisco International (SFO), San Jose (SJC), and Oakland (OAK). Another airport, Buchanan Field (CRC), has had sporadic commercial service to Southern California, but is not considered in this section due to the lack of necessary data. A map of the Bay Area with the location of these airports is shown in Figure 5. In 1991, SFO had 70 per cent of the region's commercial enplanements, while SJC and OAK each had 15 per cent. SFO's share of the domestic market is somewhat less -- approximately 60 per cent.

The market shares of the three airports vary substantially for different destinations. Figure 6 shows SFO's 1991 market share for a set of 44 different domestic origin and destination markets. Figure 6 shows a large variation in SFO's market share, with a range between 25 and 90 per cent. The figure also shows that SFO's market share does not correlate with market size, but that it does correlate somewhat with market distance. Even within the distance categories, however, there is substantial variation.

The genesis of the positive feedback model lies in our attempts to explain this variation in market share. Our initial focus was on market characteristics, such as size, distance, proportion of trips originating in the Bay Area versus the destination point, and so forth. While some factors -- notably distance, as illustrated in Figure 6 -- had statistically significant effects, the majority of the variation in market shares remains unexplained when these effects were controlled for. Further, some of the effects that were observed lacked a clear interpretation. For example, since all three airports have the facilities to support long-haul flights, there is no obvious reason why SFO would have a stronger advantage in long-haul markets.

One interpretation for this unexplained variation in airport market share is that airlines, through their service supply decisions, exert a strong exogenous influence on the distribution

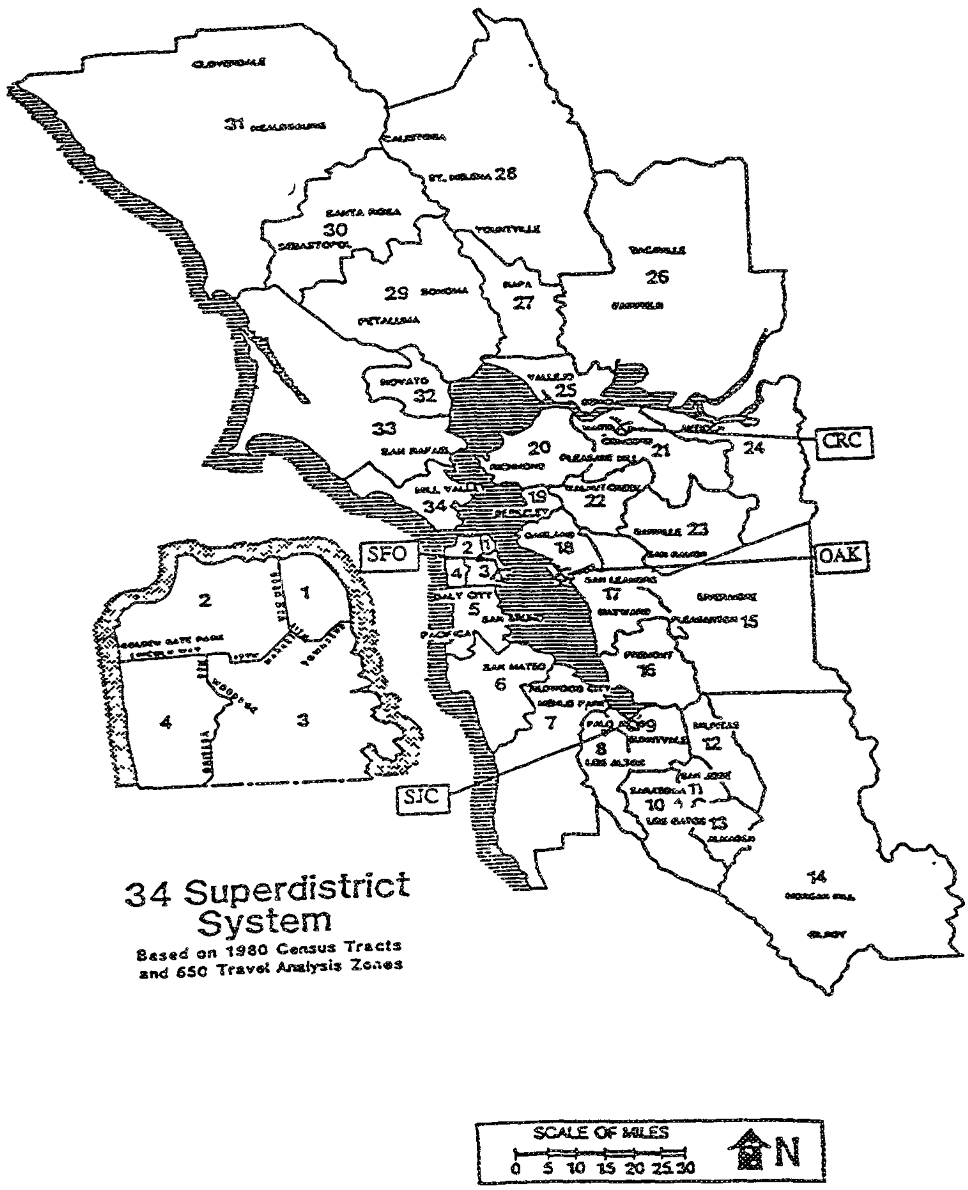

Figure 5. San Fransisco Bay Area Airports and MTC Superdistricts

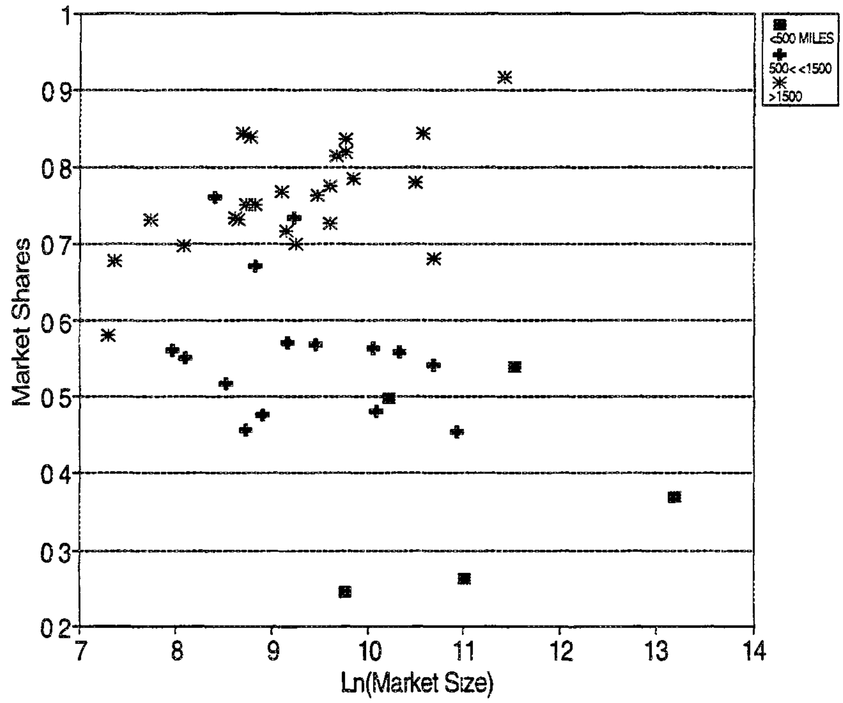

Figure 6. SFO Market Share vs. Total Market Size, 1991

of traffic in the Bay Area MAS. This impact is exogenous in the sense that it cannot be readily explained by the inherent characteristics of the markets being served. Examples of exogenous influences on airline behavior include hubbing, historical ties to particular airports, competitive strategies (de Neufville and Gellerman, 1973), and limited information.

The positive feedback model implies an alternative explanation. In this model, variation in market share of the Bay Area airports derives from variation in the relative proximity to the trip origins, which in turn derives from market-to-market variation in trip origin locations in the Bay Area. This is the same as the phenomenon as shown in Figure 4, if one imagines variation in δ resulting from differences in trip origin rather than airport location.

In applying the positive feedback model, we adopt as a working hypothesis that airline service supply, at least in the aggregate, follows demand. Since supply variables are determined by passenger traffic levels, they do not need to be introduced directly into the model. While we recognize that the working hypothesis is not literally true, our model will still have predictive value so long as a large proportion of supply behavior is endogenous. Later on, we will return to this question by relating observed supply of airline service to the passenger market share predictions obtained from our model.

4.2 Data

To apply our model, we required information concerning the origin locations of air trips between the Bay Area and different destinations, on the cost of travel between origin points and the three Bay Area airports, and on the choice of airport of tripmakers, conditional on origin location and trip destination. This section describes how we assembled these data.

We estimated the distribution of passenger origins, by destination, using the 1990 MTC Air Passenger Survey and the U.S. DOT 10 per cent airline ticket survey (as obtained from the O&D

Plus CD-ROM). The MTC survey was administered using in-person interviews of departing passengers waiting in passenger boarding areas at the three airports. 1400 flights were sampled during the seven-day survey period beginning August 13, 1991. The survey was designed to cover a representative set of flights from each airport, but the sample obtained is not representative of Bay Area air travelers as a whole, since SJC and OAK passengers are oversampled relative to SFO passengers. Sample sizes ranged from 5878 at SFO to 7895 at SJC. Table 2 summarizes and compares MTC and U.S. DOT survey sample sizes, by airport, for 44 of the larger Bay Area markets.

Surveyed passengers were asked to identify their trip origin. This information was used to identify the MTC superdistrict from which the trip originated. (MTC divides the Bay Area into 34 such superdistricts, covering the entire nine-country region; see Figure 5). Since the survey oversamples SJC and OAK passengers, it is necessary to make corrections in order to estimate the true distribution of origin points. The corrections are based on the 10 per cent survey. These data were used to compute destination- and airport-specific weights for MTC survey observations. Specifically, the weight is calculated as:

$$W_{jk} = \frac{PAXDOT_{jk}}{PAXMTC_{jk}} \tag{6}$$

where:

W_{jk} is the weight for Bay Area airport j and destination k;

$PAXDOT_{jk}$ is the 1991 passengers from airport j to destination k, as obtained from the DOT 10 per cent survey;

$PAXMTC_{jk}$ is the 1991 passengers from airport j to destination k in the MTC survey, excluding passengers whose origin in the Bay Area could not be determined or whose origin was not in the nine-county region.

Using these weights, the distribution of passenger origin points i

Table 2.
Survey Responses, by Airport and Destination

Destination	San Francisco (SFO)			Oakland (OAK)			San Jose (SJC)		
	O&D Survey	MTC Survey	Ratio	O&D Survey	MTC Survey	Ratio	O&D Survey	MTC Survey	Ratio
ANCHORAGE	2238	14	160	613	28	22	351	13	27
ALBUQUERQUE	7229	24	301	3932	66	60	1553	27	58
ATLANTA	12817	42	305	1150	20	58	1792	13	138
AUSTIN	3523	7	503	1193	13	92	2683	27	99
BOSTON	28012	58	483	1091	35	31	6843	56	122
CHICAGO	29887	181	165	5908	187	32	8109	131	62
CLEVELAND	5037	26	194	547	9	61	1117	10	112
DENVER	16944	51	332	5311	100	53	8008	159	50
DALLAS	12978	41	317	3303	54	61	6775	137	49
DETROIT	11461	37	310	1089	25	44	2231	23	97
EL PASO	2572	14	184	1986	19	105	411	5	82
MIAMI	14383	31	464	857	9	95	1966	31	63
HOUSTON	14778	55	269	2195	32	69	1838	29	63
HARTFORD	4615	14	330	566	10	57	968	9	108
LOS ANGELES	19854	980	203	97527	1620	60	68378	830	82
LAS VEGAS	13700	64	214	7847	177	44	5888	119	49
KANSAS CITY	7461	16	466	1458	47	31	1234	20	62
MINNEAPOLIS	10739	90	119	782	13	60	3260	40	82
NEW ORLEANS	6770	20	339	1104	22	50	937	11	85
NEW YORK	82697	285	290	2086	56	37	5361	91	59
OMAHA	1806	6	301	612	12	51	857	6	143
ONTARIO	15866	42	378	33156	356	93	11686	208	56
ORLANDO	9921	18	551	881	13	68	2187	25	87
PHILADELPHIA	14149	71	199	1021	22	46	2093	27	78
PITTSBURGH	5397	42	129	324	19	17	720	9	80
PHOENIX	25547	70	365	23727	179	133	6899	117	59
SEATTLE	23705	207	115	10269	285	36	9642	202	48
SAN DIEGO	54406	204	267	29357	595	49	17314	188	92
SPOKANE	3386	37	92	543	22	25	528	13	41

Destination	San Francisco (SFO)			Oakland (OAK)			San Jose (SJC)		
	O&D Survey	MTC Survey	Ratio	O&D Survey	MTC Survey	Ratio	O&D Survey	MTC Survey	Ratio
SAN ANTONIO	4547	17	267	1401	15	93	834	7	119
ST LOUIS	7179	42	171	1070	18	59	1994	32	62
TAMPA	4120	16	258	539	11	49	978	9	109
TUCSON	2805	9	312	1536	23	67	1802	18	100
DC	32754	68	482	1851	43	43	4212	35	120
ALBANY	1676	7	239	194	7	28	419	2	210
BALTIMORE	6703	10	670	940	12	78	1719	11	156
GRAND RAPIDS	1070	7	153	147	6	25	358	5	72
INDIANAPOLIS	4014	12	335	451	19	24	1007	20	50
LITTLE ROCK	856	4	214	340	7	49	281	4	70
MILWAUKEE	4972	19	262	252	6	42	681	9	76
PORTLAND	11480	141	81	5272	203	26	7142	87	82
RENO	4301	15	287	10275	6	1713	2899	37	78
SALT LAKE	5467	26	210	2028	34	60	2061	48	43
TULSA	1613	9	179	694	5	139	568	13	44

over the MTC superdistricts was estimated as follows:

$$PAX_{i\,k} = \sum_{j} PAXMTC_{ijk} \cdot W_{jk} \tag{7}$$

This procedure cannot be used when $PAXMTC_{jk}$ is zero, and is subject to high errors when $PAXMTC_{jk}$ is small. To establish some minimal level of accuracy, we eliminated from consideration all destinations that did not have at least 10 MTC Survey passengers for each airport. This leaves a total of 44 destinations -- the ones included in Table 2 -- for use in the subsequent analysis.

Figure 7 illustrates the county distribution of trip origins for several representative markets. As hypothesized, there is marked variation in the distributions. For example, San Francisco county originates the plurality of trips to Los Angeles and San Diego, while Sana Clara is the leading generator of Boston and Phoenix traffic. Among the smaller counties, Sonoma accounts for just 2 per cent of Boston passengers, but 6 per cent of those going to Ontario. The comparisons support the conclusion that variation in trip origin locations account for much of the variation in airport market share.

The final data item required for the model is a measure of travel cost between each MTC superdistrict and each airport. For this purpose, we used AM Peak period drive-alone travel times skimmed for MTC's transportation network. These travel times are defined for MTCs 700 traffic analysis zones. We estimated times at the superdistrict level to be those from the traffic analysis zones containing the superdistrict centroids. Travel times are used as a proxy for all costs associated with travel to the airport. Of course, not all airport trips are made by auto, and not all are made in the AM Peak. We argue below that systematic differences among airports in the relation of the chosen travel cost metric to the "true" travel cost may account for differences in the airport specific utilities estimated in the airport choice model.

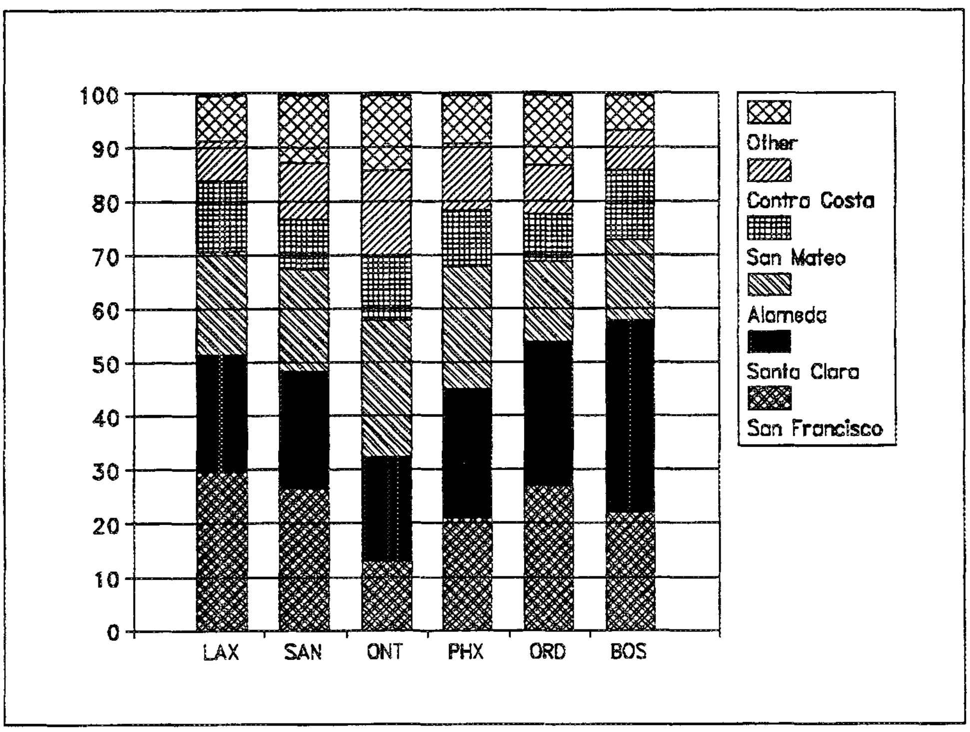

Figure 7: Air Traffic Market Shares of Bay Area Counties

4.3 Airport Choice Model

The first analytical step in the modelling process was to estimate an airport choice model. The MTC Air Passenger Survey was used for this purpose. It was again necessary to adjust the survey results to reflect the actual airport traffic levels, by destination. Each individual passenger observation was therefore weighted by a factor proportional to W_{jk}, but normalized so that the sum of the weights equaled the number of usable observations--10,500.

The utility function in Equation 2 is not linear in the parameter θ. Although the logit estimation software (ALOGIT) has the capability to estimate coefficients of this type, efforts to do so proved unsuccessful. This may be because the term θ multiplies is essentially constant for each airport (actually, it varies slightly since it includes all passengers other than those going to destination k). Thus, the data set is ill-conditioned from the standpoint of estimating θ. To overcome this problem, we performed the analysis for several assumed θ values, ranging from 0 to 10^{-1}. At θ=0, only market traffic influences the first term in the utility function, while at $\theta=10^{-1}$, this term is dominated by airport traffic except in the largest markets.

Table 3 summarizes logit estimation results. All terms are of expected sign and statistically significant at the .05 level in all models, with the exception of the distance term when $\theta=10^{-4}$. The performance of the models is fairly uniform, although there is some fall-off in the log likelihood as θ increases beyond 10^{-3}. The travel time coefficient estimate is very stable across all models. The airport constants are also fairly stable. It appears that, ceteris paribus, OAK is most attractive to passengers, and SJC is least so. Based on the ratio of the airport constants to the access time coefficient, the advantage of OAK is equivalent to a 3-5 minute reduction in access time, while the disadvantage of SJC is equivalent to a 4-5 minute increase in this time, relative to SFO. As suggested above, this could result from systematic differences in the relation between the access time measure used in the model and the actual cost of accessing the airport. Other possible explanations for these differences in airport attractiveness include inter-airport differences in airline supply behavior, congestion levels (Harvey, 1987), and amenity levels.

The value of θ has a pronounced effect on the coefficient estimates in the two traffic terms in the model. When $\theta=0$, 10^{-4}, and 10^{-3}, α, the coefficient on the first traffic term (that involving θ) in (1) is high, and increases with θ. On the other hand, for these θ values, β, the coefficient on the second traffic term (that including distance), is low. For larger values of θ, α decreases, while β increases.

Despite the substantial variation in coefficient estimates, the various models lead to very similar relationships between airport utility and market traffic, except when market traffic is small. To show this, in Figure 8 we plot the traffic component of the utility (that is, the sum of the two traffic terms in the utility function), relative to a market traffic level of 50 thousand passengers, against market traffic for the models with $\theta=0$, 10^{-3}, and 10^{-1}. When traffic is less than 10 thousand per year, the models with higher θ yield a higher utility, but this difference becomes negligible at higher traffic levels. Thus, as intended, the effect of the θ term is to increase an airport's attractiveness when traffic is too low to support non-stop service. Interestingly, however, the $\theta=10^{-3}$ model yields higher utility at low traffic levels than the $\theta=10^{-1}$ model. This reflects the sharp reduction α when θ is increased from 10^{-3} to 10^{-1}.

In terms of the log likelihood function, the model with $\theta=10^{-3}$ performs best. The performance of the model, however, should be judged according to the accuracy of the equilibrium market share predictions it yields as well as the log likelihood. We will see below that other θ values yield more accurate market share predictions, and that consequently we cannot unambiguously identify the "best" value for θ.

The models presented in Table 3 are attribute-specific, reflecting the maintained hypothesis that the attractiveness of

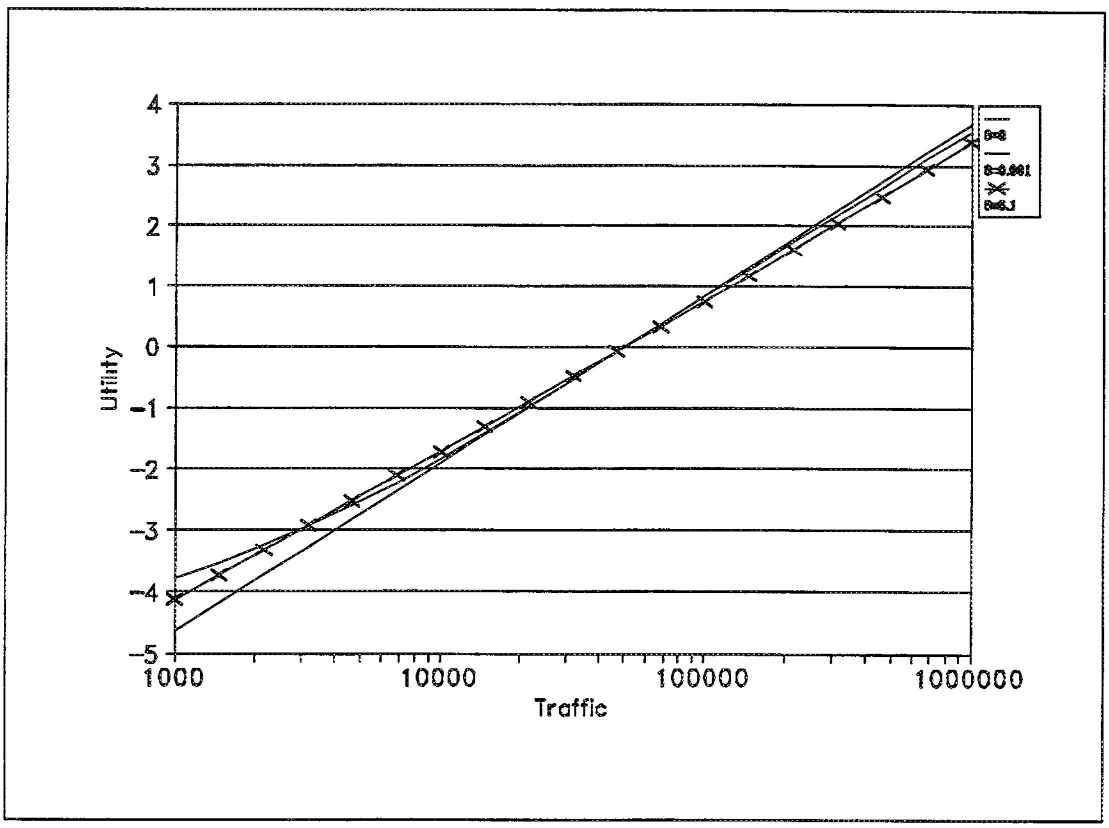

Figure 8. Airport Utility vs Traffic, Distance=1000 Miles, Enplanements=3 Million

each airport is similarly affected by changes in the traffic and access time variables. In order to consider this hypothesis, we also estimated the model in alternative-specific form in which the coefficients in the utility function are allowed to vary among the three airports. A comparison of results, for the $\theta=10^{-2}$ model, is presented in Table 4. The coefficient estimates in the alternative specific model are quite consistent with those in the attribute specific model, with deviations between estimates within standard errors with the exception of the SJC constant. Since the arguments underlying the model presuppose general, airport invariant, relationships, this consistency lends some support to them. The hypothesis that the attribute-specific model explains airport choice as well as the alternative-specific one can be tested formally, since twice the difference in the log likelihood functions is χ^2 distributed with degrees of freedom equal to the number of additional parameters in the alternative-specific model (Ben-Akiva and Lerman, 1985). This test indicates that, despite the similarity of estimates, the null hypothesis that the coefficients are actually equal can be rejected with a high (better than 99 per cent) level of confidence. Nonetheless, since the attribute specific model is simpler, more consistent with the underlying theory, and has coefficients that are for the most part quite close to those in the alternative-specific model, we will employ them in the subsequent analysis.

4.4 Calculation of Equilibria and Analysis of Model Performance

The logit model results, combined with the information on the distribution of passenger origin locations, were used to calculate airport traffic equilibria. The approach is similar to that outlined in Section 2 (Equations 1-3), except that for each individual market calculation, the <u>observed</u> airport traffic to other markets, rather than the traffic predicted by the model was used. In adopting this approach, we are in effect testing the performance of the model at the market, rather than the airport level. This avoids the problem of error propagation arising when

Table 4. Logit Estimation Results,
Attribute-Specific and Alternative-Specific Models

Model

$$V_{ijk} = \alpha \cdot \log(PAX_{jk} + .01 \cdot (NLPAX_j + \sum_{n \neq j} PAX_{jn})) +$$
$$\beta \cdot \log(DIST_k) \cdot \log(1 + PAX_{jk}) + \psi \cdot ATIME_{ij} + \gamma_j$$

Coefficient	Attribute-Specific Model			Alternative-Specific Model		
	SFO	OAK	SJC	SFO	OAK	SJC
α	0 4692 (.188)	0 4692 (188)	0 4692 (188)	0 5120 (215)	0 4680 (206)	0 4623 (218)
β	0 1187 (016)	0 1187 (016)	0 1187 (016)	0 1233 (017)	0 1288 (021)	0 1355 (020)
ψ ($\min^{-1}$)	-0 1027 (002)	-0 1027 (002)	-0 1027 (002)	-0 1072 (019)	-0 1122 (023)	-0 0932 (020)
γ_j	0	0 4750 (097)	- 3662 (126)	0	9211 (480)	-1 0290 (486)
log likelihood	-4945 5			-4918 1		
ρ^2	4982			5010		

Standard errors in parantheses

inaccurate predictions in one market, by affecting airport traffic, cause inaccurate predictions in another market. Moreover, this approach allows separate calculations of equilibria by market, substantially reducing the computation burden involved.

To calculate the equilibria, we took the actual airport market shares as a starting point, and calculated new shares using Equations 1 and 2. The new shares were then used in a second airport market share calculation, and iteration continued until stable market shares were obtained. This procedure is much the same as that depicted by arrows in Figure 1, and will thus naturally arrive at stable, rather than unstable, equilibria. If there are multiple stable equilibria, however, the procedure will tend to identify those which are "closest" to the observed market shares. We have not thoroughly explored the possibility of multiple stable equilibria in a three-airport MAS, but did conduct several numerical experiments using different (non-zero) starting values for market share, and always arrived at the same equilibrium result.

Figures 9-16 compare the equilibrium traffic share to the observed traffic share (based on the DOT 10 per cent survey). Two plots are presented for each value of θ up to 10^{-2}. In the first plot, results for all 44 markets are included. In the second plot, only the 13 largest markets (based on total traffic from all three airports) are shown. The equilibrium and observed values are expected to be closer for the large markets for two reasons. First, these market are better represented in the MTC survey, and thus the distribution of passenger origin locations can be estimated with greater accuracy. Second, since these markets have more responses in the DOT survey, airport market shares are also observed with greater accuracy.

Inspection of the plots confirms this expectation. Predictions for large markets correlate quite well with observation, while the fit is much worse when all markets are considered. Most importantly, the large market plots reveal strong correlation of predicted and actual markets shares for the individual airports,

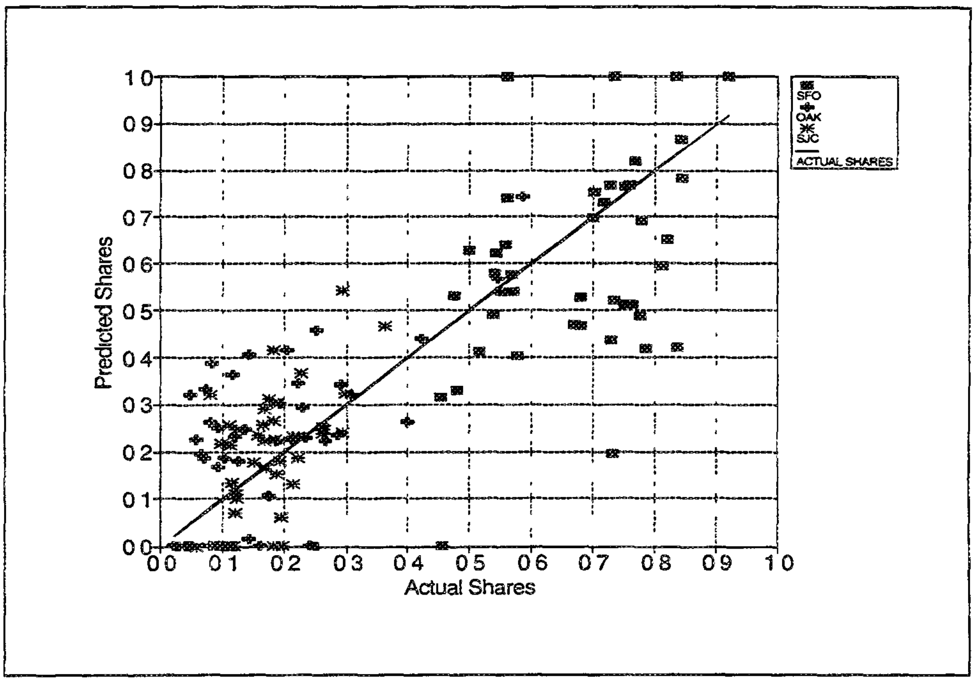

Figure 9. Predicted vs. Actual Market Shares, θ=0, All Markets

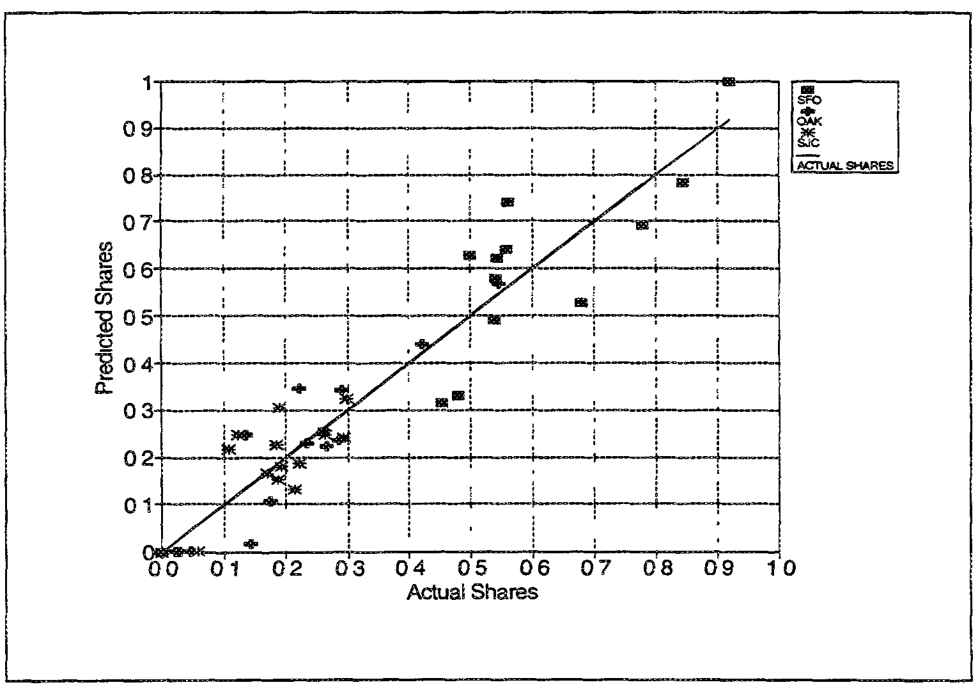

Figure 10. Predicted vs. Actual Market Shares, θ=0, Large Markets

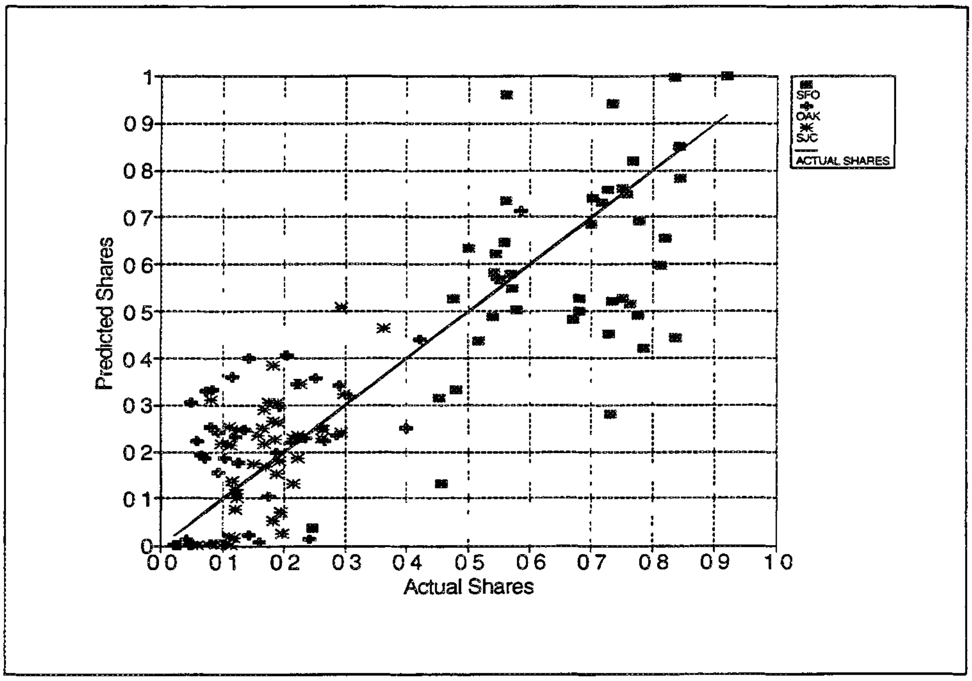

Figure 11. Predicted vs. Actual Market Shares, θ=0.0001, All Markets

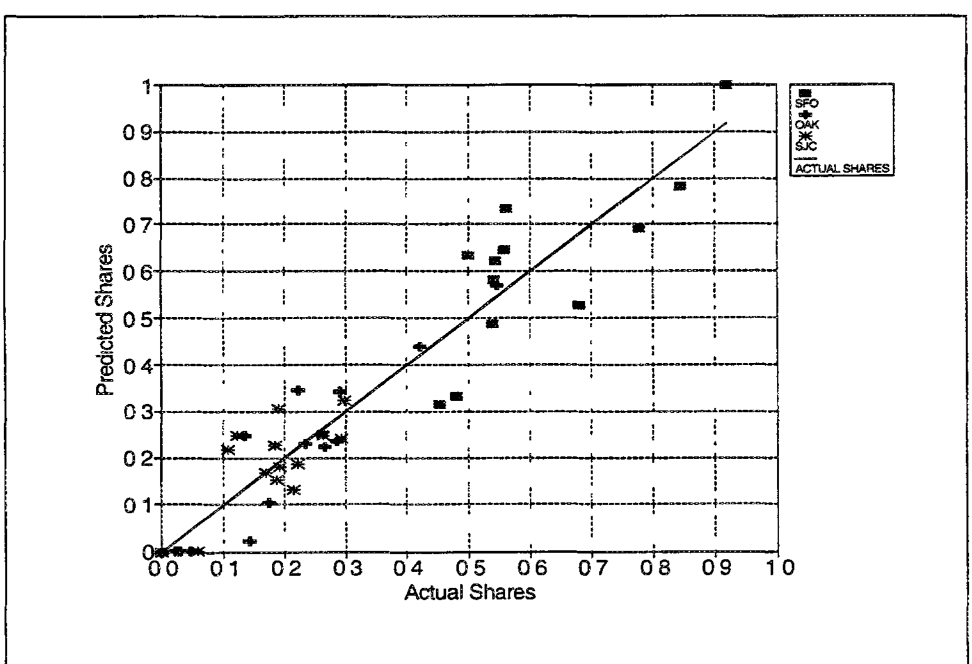

Figure 12. Predicted vs. Actual Market Shares, θ=0.0001, Large Markets

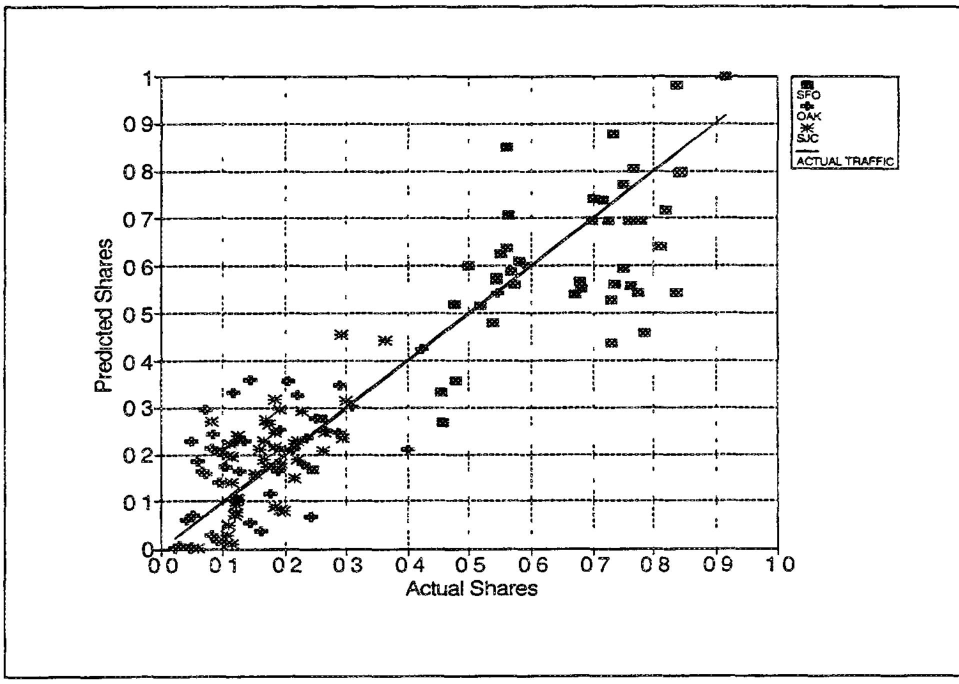

Figure 13. Predicted vs. Actual Market Shares, θ=0.001, All Markets

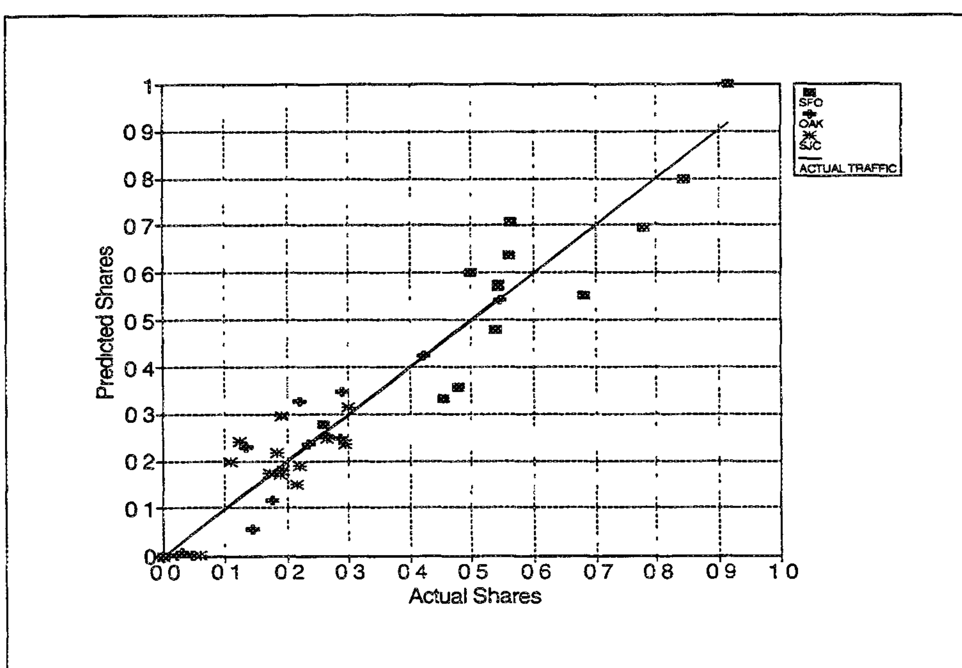

Figure 14. Predicted vs. Actual Market Shares, θ=0.001, Large Markets

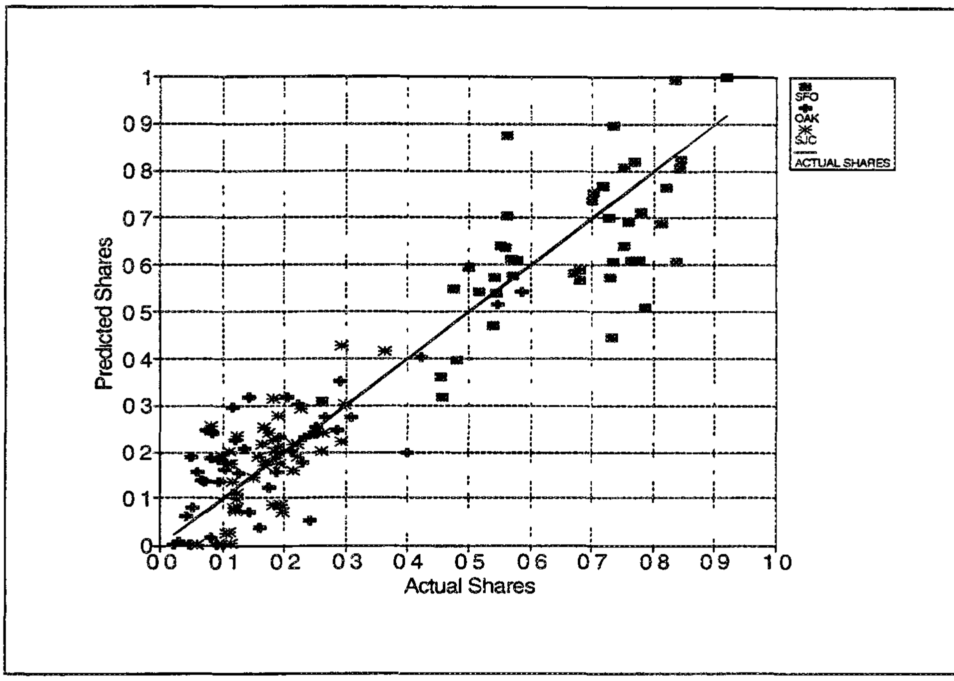

Figure 15. Predicted vs. Actual Market Shares, θ=0.01, All Markets

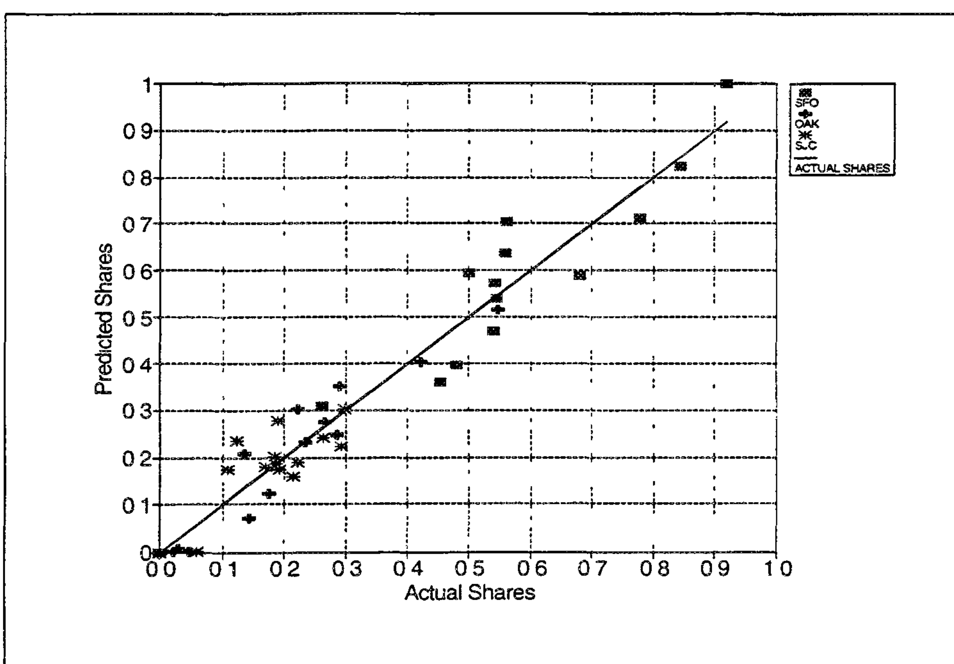

Figure 16. Predicted vs. Actual Market Shares, θ=0.01, Large Markets

not just for the set of observations as a whole. This is crucial, since by using airport constants in the logit utility functions, we have essentially guaranteed that our model will correctly "predict" mean market shares of the three airports. The additional predictive power of the model used here rests in its ability to explain market-to-market variation around the mean values.

Comparing the plots for all markets for the different θ values, it is apparent that the models with $\theta=10^{-3}$ and 10^{-2} have a better fit than those with $\theta=0$ and 10^{-4}. The former models are more likely to predict equilibria with market shares between 0 and 1, while the latter tend toward outcomes in which airports receive all or none of the traffic. These results are consistent with Figure 4, which shows that the models with $\theta=10^{-3}$ and 10^{-2} predict higher utility levels for airports with low traffic than those with θ values below or above this range. These higher utility levels allow airports with low traffic levels in a given market to retain some market share. Since these differences are confined mainly to low traffic situations, model performance is much less variable when only large markets are considered.

It is instructive to compare the plots for the various positive feedback models with the results for a model in which only access time (in addition to airport constants) is considered -- that is, a model in which α and β in Equation 2 are assumed to be 0. Figures 17 and 18 present the plots for this case (the plots for $\theta=10^{-1}$ look similar). It is clear that this model is inferior to the models incorporating positive feedback. The former predicts a much smaller amount of market share variation than what is observed and what the positive feedback models predict. This is precisely what we should expect: the positive feedback effect amplifies slight market-to-market variation in airport accessibility into much larger differences in airport market share.

To more formally compare and assess the performance of the models, we performed statistical tests of hypotheses pertaining to their ability to explain variation in airport market shares. For reasons previously stated, we concern ourselves only with intra-

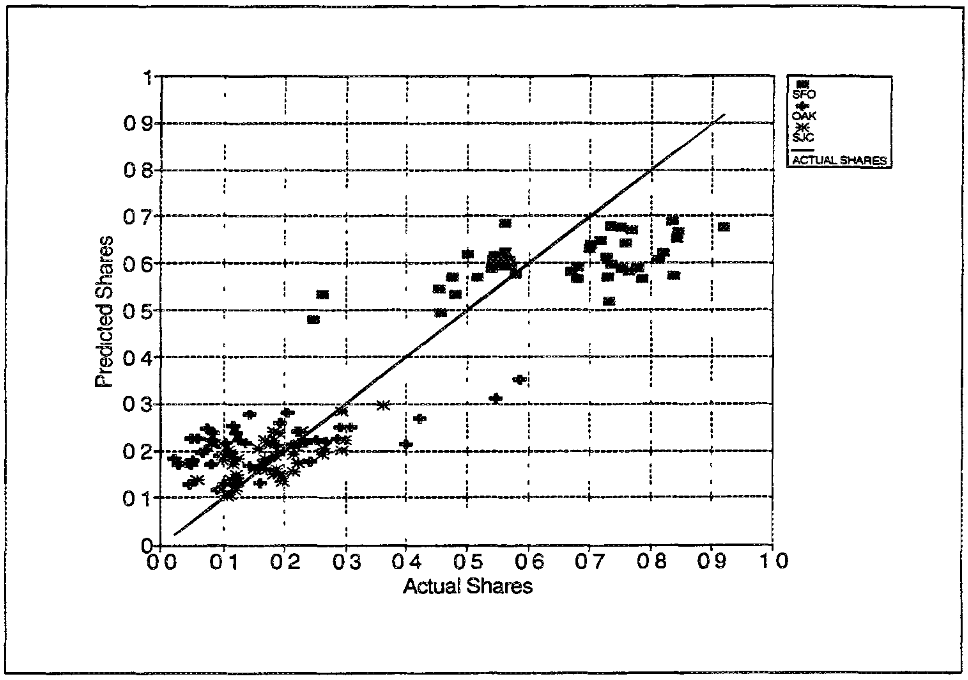

Figure 17. Predicted vs. Actual Market Shares with only Travel Time Considered, All Markets

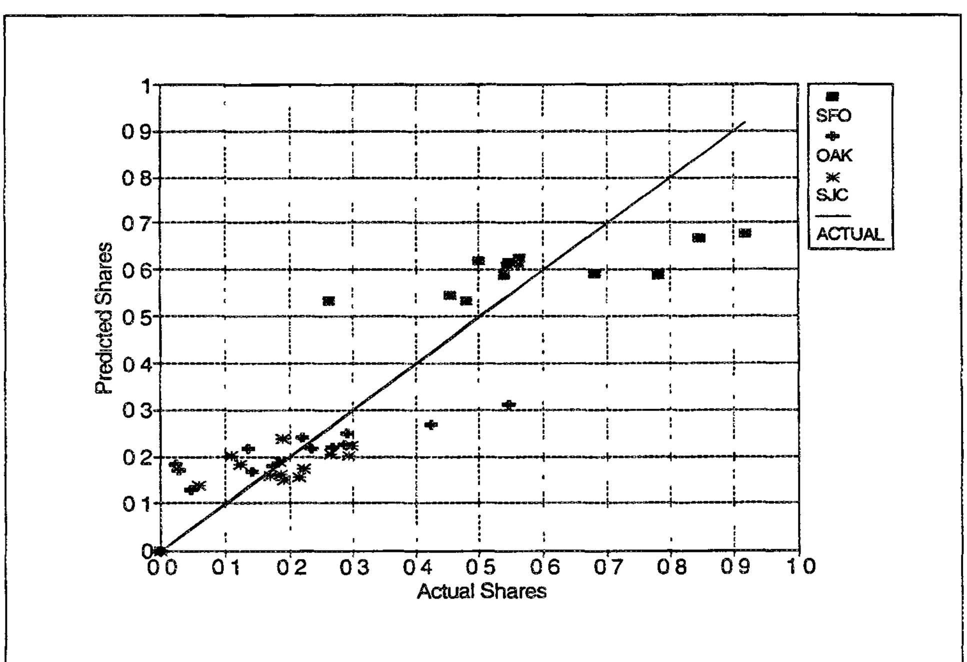

Figure 18. Predicted vs. Actual Market Shares with only Travel Time Considered, Large Markets

airport variation.

Two different null hypotheses can be used to assess model performance. The first, which we term "null hypothesis 1," is that our model predicts the market share of airport j in destination k, MS_{jk}, no more accurately than the model:

$$MS_{jk} = \overline{MS}_j \qquad \textbf{(8)}$$

where:

$\overline{MS}_j$ is the mean airport share of airport j over all markets.

An F-statistic can be used to test null hypothesis 1. The test statistic is given by:

$$f(n-1,n-1) = \frac{\sum_k (MS_{jk} - \overline{MS}_j)^2/n-1}{\sum_k (MS_{jk} - \hat{MS}^m_{jk})^2/n-1} \qquad \textbf{(9)}$$

where:

$\hat{MS}^m_{jk}$ is the market share for airport j to destination k predicted by model m.

n is the number of markets considered in the analysis (n=44).

Tables 5 and 6 summarize results for each airport, and for each of the five models. Table 5 is based on all markets, and Table 6 on the 13 largest. The better performance of the models for large markets is again evident. It is also apparent that the model performance varies according to the airport, with OAK predictions the best, followed by SFO, and finally SJC. For the former two airports, and for all θ values, null hypothesis 1 can be rejected at a level of significance of 5 per cent or better. This null hypothesis cannot be rejected for large markets in the case of SJC, however. Models with larger θ values tend to perform somewhat better for large markets, although no single model performs best

Table 5.
F-Statistics, Null Hypothesis 1, All Markets

Airport	θ=0	θ=0.0001	θ=0.001	θ=0.01	θ=0.1
SJC	0.39	0.44	0.68	0.81	0.51
SFO	0.67	0.75	1.20	1.57*	0.99
OAK	1.06	1.14	1.60**	2.05***	1.50*

*Significant at 10 per cent level
**Significant at 5 per cent level
***Significant at 1 per cent level

Table 6.
F-Statistics, Null Hypothesis 1, Large Markets

Airport	θ=0	θ=0.0001	θ=0.001	θ=0.01	θ=0.1
SJC	1.26	1.25	1.23	1.51	1.74
SFO	3.66**	3.61**	3.55**	4.73***	5.46*
OAK	6.71***	6.70***	7.07***	9.47***	9.02***

*Significant at 10 per cent level
**Significant at 5 per cent level
***Significant at 1 per cent level

for all three airports. Results are less convincing when all markets rather than just the largest ones are considered. Null hypothesis 1 can be rejected at a 5 per cent significance level only for the model with $\theta=10^{-2}$ and only for OAK. This model yields the best results for all three airports when all markets are considered.

A second null hypothesis, "null hypothesis 2," is that there is no linear relationship between an airport's market share and the share predicted by our model. This translates to the hypothesis that in the linear model,

$$MS_{jk} = \beta_j^m + \pi_j^m \cdot \hat{MS}_{jk}^m + \epsilon_{jk}^m,$$

$\pi_j^m = 0$. Null hypothesis 2 is rejected if there is a statistically significant (and positive) linear relationship between the predicted and observed market shares. Such a relationship can exist even when differences between predicted and observed values, upon which the testing of null hypothesis 1 is based, are large.

Null hypothesis 2 was tested using linear regression. Tables 7 and 8 summarize the regression results. As before, the analysis is conducted for each airport and each model, using both the entire set of 44 markets and the 13 largest markets. In general, these results show that the null hypothesis must be rejected at the 1 per cent significance level, for all models and airports, and whether all markets of large markets are considered. Thus, even when equilibrium market share is not itself an accurate predictor of actual market share, there remains a strong linear relationship between these variables.

The regression results display many of the same patterns as the F-statistics discussed above. When all markets are considered, the best results are for models with $\theta=10^{-2}$, but when large markets are considered, the model with $\theta=10^{-1}$ yields a somewhat higher R^2. Like the F-statistics, the regression results show that market shares in large markets can be predicted more accurately, and that OAK predictions are the most accurate, followed by SFO, with

Table 7.
Regression Results, All Markets

Airport	Coefficient	θ=0	θ=10^{-4}	θ=10^{-3}	θ=10^{-2}	θ=10^{-1}
OAK	INTERCEPT	0 053	0 050	0 034	0 030	0 043
	SLOPE	0 554 (093)	0 575 (094)	0 685 (097)	0 760 (0 100)	0 657 (0 101)
	R^2	0 456	0 468	0 541	0 528	0 500
SFO	INTERCEPT	0 407	0 388	0 291	0 239	0 338
	SLOPE	0 417 ((088)	0 448 (090)	0 598 (0 097)	0 657 (096)	0 507 (091)
	R^2	0.353	0 370	0 474	0 528	0 425
SJC	INTERCEPT	0 117	0 112	0 094	0 092	0 112
	SLOPE	0 297 (066)	0 319 (069)	0 416 (077)	0 450 (079)	0 346 (070)
	R^2	0 324	0 340	0 410	0 437	0 366

Standard errors in parantheses

Table 8.
Regression Results, Large Markets

Airport	Coefficient	θ=0	θ=10^{-4}	θ=10^{-3}	θ=10^{-2}	θ=10^{-1}
OAK	INTERCEPT	0 037	0 037	0 037	0 030	0 025
	SLOPE	0 841 (092)	0.838 (092)	0 839 (.088)	0 890 (084)	0 916 (.092)
	R^2	0 883	0 884	0 892	0 910	0 900
SFO	INTERCEPT	0 128	0.130	0 132	0 096	0 080
	SLOPE	0 791 (126)	0.788 (126)	0 784 (126)	0 833 (117)	0 857 (112)
	R^2	0 783	0 781	0 778	0 822	0 842
SJC	INTERCEPT	0.074	0 075	0 076	0 064	0 056
	SLOPE	0 590 (205)	0.587 (205)	0.581 (.207)	0 656 (211)	0 704 (.204)
	R^2	0 431	0 428	0 418	0 469	0 518

Standard errors in parantheses

results for SJC the least satisfactory. On the other hand, whereas F-statistics indicated that equilibrium market share is not itself a good direct predictor of SJC market shares, the regression results show that a linear function of the equilibrium market share is a good predictor, explaining close to half the variation in the observed value.

All regressions yield positive intercepts and slopes less than 1. This pattern admits of both a substantive and a statistical explanation. The substantive interpretation is that positive feedback effects are somewhat damped by other factors not accounted for in our model. We have already discussed in general terms how airport market shares are influenced by "exogenous" airline actions which we have intentionally omitted. Such influence would be expected to attenuate the relationship between actual market shares and those predicted by our model in the manner the regression results suggest. Alternatively, the regressions estimates could be distorted by error-in-variables bias. Specifically, calculated equilibria are subject to error because they are based on estimates of the distribution of trip origins that are subject to error. When there is significant measurement error in an independent variable used in a regression, it is expected that the estimated coefficient on the variable will be of smaller absolute value than the true value of the coefficient. This is the probable explanation for the slope estimates being greater in the large market analyses, where a greater number of MTC responses results in a smaller error in the estimated distribution of trip origins.

Before leaving this section, we consider the question of which θ value is "best". Different criteria favor different values. The best logit results are obtained for the smaller θ values -- 0, 10^{-4}, and 10^{-3}. Higher θ values -- 10^{-2} and 10^{-1} -- yield models that predict market shares more accurately. The probable reason for this conflict is related to the error-in-variables problem. Although models with smaller θ predict airport choice somewhat better, models with larger θ may be more robust with respect to errors in trip origin distribution. In any case, application of the model to a single MAS does not provide an adequate basis to determine θ, or even to decide whether the θ term is appropriately specified. However, we have chosen the $\theta=10^{-2}$ model for use in the subsequent sections of this paper. Since this model is arguably the best in terms of market share prediction, and only slightly below the best as measured by logit model performance, it seems the most reasonable compromise. All the other models are decidedly inferior with respect to one or the other of these criteria.

5. Equilibrium Market Share and the Supply of Air Service

As discussed in Section 2, our model is based on the maintained hypothesis that the supply of air service is endogenous -- determined by traffic rather than vice versa. The high correlation between predicted and observed market shares found in the last section lends some credibility to this assumption. In this section, we examine this matter further by considering relationships between service supply and predicted market share.

Our measure of service supply is available non-stop seats, as obtained from the USDOT service segment data base. We consider the relationship between an airport's share of non-stop seats and market share in the 13 largest Bay Area markets. In general, supply of non-stop seats in these markets was well in excess of origin-and-destination passengers, reflecting the presence of considerable non-local traffic on these segments. This in turn is a result of hubbing, both through the Bay Area airports SJC and SFO, and at the destination airport. Overall there were between 2 and 3 times as many seats as local passengers in these markets, with SJC having the highest ratio, followed by SFO. OAK, the only Bay Area airport not used as a hub, had the lowest ratio.

Figure 19 plots observed seat share against the equilibrium market share (based on the model with $\theta=10^{-2}$). The plot confirms that these variables are correlated, both overall and at the airport level. The relationship between seat share and market share is observed to differ according to the airport. SFO shares are highly correlated, with the seat share generally higher than the market share. OAK shares are also highly correlated, but in this case the market share is typically higher than the seat share. SJC displays noticeably weaker correlation than the other airports. Like SFO, its seat share generally exceeds its market share. Regression results, summarized in Table 9, confirm these impressions. It is clear from these results that for a given equilibrium market share, OAK has a lower expected seat share than either SFO or SJC.

These results closely parallel those in the last section,

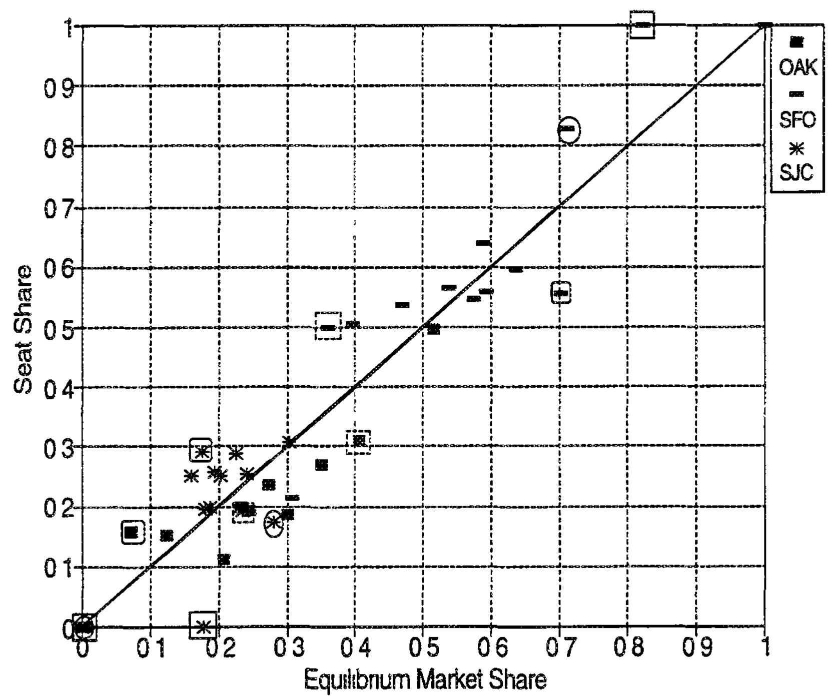

Figure 19. Observed Seat Shares vs Equilibrium Market Share, Bay Area Airports

Table 9.
Regression Results, Seat Share Models

Coefficient	Linear Model	Log-Linear Model
Intercept	-0 315	0 100
Sfdummy	0 133 (0 165)	0 084 (0 051)
Oadummy	-0 192 (0 116)	-0 043 (0.034)
Predicted Traffic Share	0 740 (0 146)	0 660 (0 134)
R^2	0 71	0 75

Standard errors in parantheses

where it was found that observed market share is more highly correlated with equilibrium market share in the case of OAK and SFO. This suggests that the weaker predictive performance of the positive feedback model in the case of SJC stems for the failure of the supply side of the system to respond to local demand. One possible reason for this may be that in the 1991 time period we are analyzing American Airline's SJC hub was fairly new, and the airline has thus not had sufficient time to tailor its schedule to local market demand. Alternatively, it may be that American's service supply decisions at SJC were dominated by connecting market considerations, substantially weakening the tie between local traffic and service. In any case, it is interesting that the San Jose hub has come to be seen as a failure, and operations there have been sharply curtailed at the time of this writing.

Figure 19 reveals some markets in which the market and seat shares diverge significantly. While we cannot offer definitive explanations of these divergences, consideration of a few specific cases may prove illuminating. Data points corresponding to four such divergent markets are identified in Figure 19. Solid rectangles indicate the Washington, D.C. market, in which SFO has 100 per cent of the seats, but a predicted market share of 82 per cent. This is an example of a long-haul market in which non-stop service may not be economically viable from SJC or OAK. Passengers using these airports to travel to Washington would be forced to rely on multi-stop or connecting service. A somewhat similar story holds for the Boston market, data for which is indicated by a solid circle in Figure 19. In this case, OAK, which is expected to garner less than 1 per cent of the market, could not support service, but SJC, for which a 28 per cent market share is predicted, does have service. SJC has a smaller share of non-stop seats than traffic, however, for two possible reasons. First, the amount of non-stop service is such that a sizable fraction of SJC passengers to Boston opt for other services. Second, SFO may be dominating the non-local passenger traffic in this market.

In the third example, Dallas (indicated by an oval), SFO has a smaller share of seats (55 per cent) than of equilibrium passengers (70 per cent), with the reverse holding for both SJC and OAK. The probable explanation for this is that many passengers out of the OAK and SJC hub through DFW on their way to destinations in the south and east which to which SFO has non-stop service (for example, Washington, as discussed above). DFW would be a particularly attractive hub for travelers out of SJC, since American has hubs at both locations.

Finally, Phoenix, indicated by a dashed rectangle, is another case in which SFO has a higher seat share (50 per cent) then predicted market share (36 per cent). The apparent explanation for this disparity is that United used its SFO hub to connect Phoenix with cities in the Pacific Northwest such as Portland and Seattle, while little connecting traffic on the America West and Southwest flights out of SFO and SJC.

The above remarks illustrate how hub-and-spoke networks and traffic thresholds for economically viable services create disparities between airline seat shares and the market shares predicted by our model. These disparities in turn result in differences between observed and predicted market shares. The suggested explanations are admittedly speculative, and further analysis is required to confirm them. The more important result is that the factors suggested in this discussion play a comparatively minor role, compared to those included in the positive feedback model, in shaping the distribution of traffic in MAS.

6. Predictions for a Fourth Airport

We used the model to forecast equilibrium market share for Buchanan Field (CRC), an airport in Central Contra Costa County, about 30 miles from San Francisco. The airport has received sporadic commercial service in the past. The enplanements at Buchanan peaked at 59 thousand in 1988. Subsequently the two airlines serving Buchanan, PSA and USAir, merged, and USAir then retrenched from the Bay Area market. These events virtually eliminated commercial activity at Buchanan. In 1991, total enplanements numbered just 5,600.

We calculated Buchanan's equilibrium market share for two destinations: Los Angeles and San Diego. We treated each market in isolation--as though it were the only one served from this airport. Thus each market share is calculated assuming that there is no traffic in any other market using Buchanan. Since we lacked an airport-specific utility term for Buchanan, we carried out the analysis parametrically, allowing this utility value to range from -0.5 (the SJC value) to +0.4 (the OAK value).

The results are shown in Figure 20. Equilibrium market shares for Buchanan range between 15 and 26 per cent for San Diego, and between 11 and 17 per cent for Los Angeles. This substantial market penetration reflects the fact that Buchanan is advantageously located. Of the four airports being considered, Buchanan is the closest to the North Bay and much of Contra Cost County. Additionally, it is fairly far from SFO, and thus offers a pronounced travel time advantage over that facility to a sizable fraction of the Bay Area market.

But while the large predicted market shares for Buchanan are reasonable, they are not born out by experience. Even in the peak year of 1987, CRC captured only 3.3 per cent of the Los Angeles market and less than 1 per cent of San Diego traffic. What happened? One explanation is that Buchanan never reached the point where it was considered as a potential alternative by most travelers. Whereas virtually any East Bay traveler will consider OAK, and any South Bay traveler will consider SJC, when making

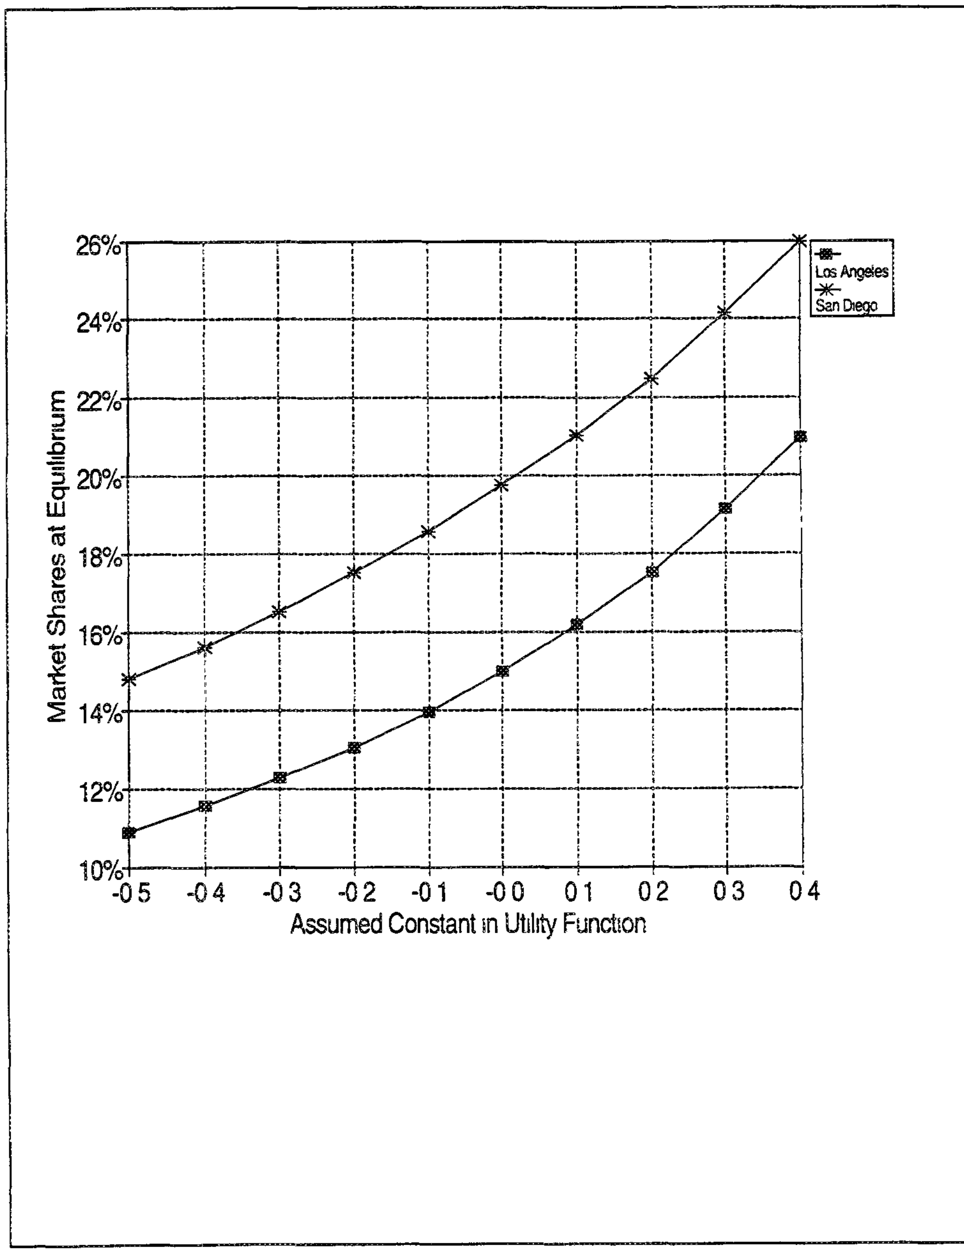

Figure 20. Predicted Market Share vs. Airport-Specific Utility, Buchanan Field

travel plans, we suspect that the possibility of using CRC never crossed the minds of most travelers in its market area. Indeed, a survey conducted by CRC revealed that many respondents did not know of the availability of commercial service there (White, 1993).

A second reason was inadequate airport infrastructure. In the late 1980s, the only commercial jet aircraft that could operate from CRC's 4600 foot runway was the BAE 146. While assessments of the mechanical reliability and performance of this aircraft vary, it was not widely prevalent in the fleets of U.S. domestic carriers, and this may have discouraged more widespread entry. Furthermore, parking was limited, and there was no permanent terminal building.

Finally, these was strong political opposition to further commercial development at CRC. The CRC Board imposed tight limitations on the number of flights, and generally discouraged airlines from expanding commercial service. This political opposition is one reason why steps were not taken to remedy the infrastructural deficiencies.

Thus, it appears that a combination of lack of passenger awareness, inadequate infrastructure, and lack of political support prevented CRC from reaching its potential as a commercial airport. Of these reasons, the first is perhaps the most significant, for it speaks to an additional type of feedback effect not captured in our model. While OAK, SFO, and SJC are all well established commercial airports of which virtually all travelers are aware, CRC never attained this status. Only when this information bias is overcome is it possible for a new airport to compete on equal terms with established competitors.

7. Conclusions

We have developed a simple model that contributes to the understanding of the allocation of passenger traffic among the three main commercial airports of the Bay Area. Our model is based on an accessibility effect, which makes airports close to passenger trip origins more attractive, and positive feedback effect, which makes airports with more traffic more attractive. We have shown that a model incorporating these effects explains much of the wide market-to-market variation in airport market shares observed in the Bay Area. The model works particularly well for large markets, for which it is possible to accurately estimate its main exogenous input, the distribution of locations from which air travelers begin their journeys. The model yields equilibrium market shares which are themselves fairly good predictors of observed values; even better predictions are obtained using linear models relating actual market share to the equilibrium share.

The most important implication of this research is that the locational distribution of trip origins plays a major role in determining airport traffic allocation. Indeed, it appears that this factor is more important than "exogenous" airline behavior. Also, at least in the Bay Area, there is substantial market-to-market variation in this locational distribution. This suggests that regional airport planners need to consider regional air travel patterns at a zone-to-destination level, and that such patterns cannot be accurately estimated by simple region-to-zone apportionment schemes. The modelling of travel patterns at the zone-to-destination level is thus an important area for future research.

The positive feedback effect amplifies the accessibility effect. From a regional perspective, SFO has a modest accessibility advantage over OAK and SJC. The access time to SFO for the average Bay Area passenger is 46 minutes, only 8 minutes less than the average to OAK and SJC. Because of the positive feedback effect, this advantage translates into a much larger market share differential. This poses a difficult challenge from a planning standpoint, since it implies that slight changes in zone-to-destination trip changes can sharply alter the distribution of traffic among airports. But it also means that transportation planners can exert considerable leverage on airport traffic distribution by taking steps to increase (or decrease) airport accessibility.

A potential objection to our analysis is that it does not sharply delineate cause and effect. This is a recurring issue in positive feedback models, which by definition see cause and effect in the same variables. It must therefore be emphasized that the positive feedback model is used to calculate an equilibrium, and the equilibrium state is determined solely by the locational distribution of trip origins. Thus the only potential source of "reverse causality" is if the locational distribution is itself endogenously determined -- for example, if more people from East Bay fly to Reno because OAK has better service to that destination. Although we cannot dismiss this possibility entirely, it seems implausible that airport access considerations would play more than a marginal role in shaping zone-to-destination travel patterns. Although often unpleasant, airport access represents a small fraction of the total cost of a typical air trip. It is far more likely that travelers consider access issues in selecting an airport than in deciding whether or where to fly. This conclusion is supported by the high incidence of cases in which passengers select an airport with less accessibility but better service.

We do not claim that our model is a complete representation of the behavior of MAS. Our results reflect this incompleteness in a number of ways. First, our logit model results and overestimates of market share for Buchanan field show the importance of airport-specific effects. From the airport constants in the logit model, we learn that, ceteris paribus, OAK is the most attractive Bay Area airport, with an advantage equivalent to a 4 minute differential over SFO, and a 9 minute differential over SJC. Further research is needed to explain why. Likewise, further work is required to explain why Buchanan field failed to reach the market potential predicted by our model. While there is interesting anecdotal evidence on this point, we cannot yet incorporate factors suggested by this evidence into a predictive model.

Comparisons of non-stop seat shares with predicted market shares show how factors related to hub-and-spoke route systems affect traffic in the Bay Area. In the case of SJC, it appears that American's hub had -- at least temporarily -- attenuated the linkage between service supply and the local traffic base. In several other instances, seat shares diverged from market shares as a result either of routing of non-local traffic through Bay Area airports, or the routing of local traffic through other hubs. It is not surprising that such effects exist. What is more remarkable is that relatively accurate predictions of Bay Area airport market shares can be made without taking them into account.

Future research should be devoted to several areas. First, the basic findings of this study should be verified through application of the model to other MAS. In addition to determining whether the positive feedback mechanism plays as important a role in other MAS as it does in the Bay Area, such application will reveal the transferability of model parameters, and shed additional light on the nature of airport-specific effects. Second, a more explicit representation of supply-side behavior should be considered as an alternative to the implicit, traffic-based, representation presented here. Such a representation should take into account threshold effects on non-stop service availability, pricing, and the impacts of hubbing. Third, as already noted, the market-to-market variation in trip origin distributions warrants further study. If the importance of such variation is as great as our results suggest, it will have serious implications for the nature and difficulty of the task of predicting airport traffic in an MAS. Finally, the model should be extended to explore how capacity limitations affect traffic distribution in an MAS, and identify how, in such circumstances, planning and policy interventions can lead to multiple airport systems that optimally balance the competing goals of accessibility, service quality, and infrastructure cost.

References

Ashford, Norman and Messaoud Benchemam (1987), Passengers' choice of airport: an application of the multinomial logit model, Transportation Research Board, 66th Annual Meeting.

Ashford, Norman (1989), Predicting the passenger's choice of airport, Airport Forum 19,3:42-44.

Augustinus, Johannes and Steve Demakopolous (1978), Air passenger distribution model for a multiterminal airport system, Transporation Research Record 673:176-180.

Bailey, Elizabeth, Stanley Graham, and Robert Kaplan (1985), Deregulating the Airlines, MIT Press, Cambridge MA.

Ben-Akiva, M. and S. Lerman (1985), Discrete Choice Analysis, MIT Press, Cambridge MA.

Gelerman, Walter and Richard de Neufville (1973), Planning for satellite airports, Transportation Engineering Journal 3:537-551.

Hansen, Mark (1990), Airline Competition in a hub-dominated environment, Transportation Research B 24B,1:27-43.

Harvey, Greig (1987), Airport choice in a multiple airport region, Transportation Research A 21A,6:439-449.

Innes, J.D., and Donald Doucet (1990), Effects of access distance and level of service on airport choice, Journal of Transportation Engineering 116,4:507-516.

Kanafani, Adib, G. Gosling, and S. Taghavi (1977), Studies in the demand for short haul air transportation, Institute of Transportation Studies, University of California, Berkeley, special report.

Lin, Feng-Bor (1977), Airport choice in a low demand region, Transportation Engineering Journal 6:711-727.

National Capital Region Transportation Board (1985), Mulitple Airport Demand Allocation Model (MADAM) Calibration and Sensitivity Analysis.

Skinner, Robert (1976), Airport choice: an empirical study, Transportation Engineering Journal 4:871-882.

White, Hal, Manager, Buchanan Field (1993), Telephone Interview, August, 1993.