CAASD's Future Air Traffic Timetable Estimator:

A Micro-Econometric Approach

Dipasis Bhadra, Jennifer Gentry, Brendan Hogan, and Michael Wells

The MITRE Corporation's Center for Advanced Aviation System Development (CAASD)

7515 Colshire Drive, McLean, Virginia 22102

INTRODUCTION

The aviation community has developed numerous tools for simulating the operational flows of the National Airspace System (NAS)[1] [1–3]. Some of these modeling capabilities are quite detailed in approximating the metrics they set out to depict. For example, CAASD's Detailed Policy Assessment Tool (DPAT) measures queuing delays occurring in the NAS throughout the various phases of flight. Taken together, these delays can reach significant levels on a bad weather day. Alternatively, other models have been developed that simulate airline schedule evolution to mitigate the effects of congestion. For instance the National Aeronautics and Space Administration/Logistic Management Institute's (NASA/LMI) model provides airlines with a series of actions they can take in response to congestion, including depeaking, off-hours operations, use of secondary airports, and using larger aircraft [3].

In all cases, these tools must be fed a schedule of flights, or timetable, which is generated outside of the model. A typical timetable contains columns for the origin airport, destination airport, departure time, arrival time, equipment type, and possibly the carrier. Note, we use the term "timetable" and not "schedule" in order to avoid possible confusion. Itineraries for unscheduled traffic, combined with that of scheduled commercial traffic, yields a timetable. Table 1 provides an example.

Table 1. **Sample Timetable**

DEP APT (step 2)	ARR APT (step 2)	EQUIP (step 3)	DEP TIME (step 4)	ARR (step 4)
JFK	BOS	AEST	10:30	11:30
JFK	IAD	CRJ1	10:15	11:31
JFK	KIN	A343	10:15	14:05
JFK	LAX	B762	10:10	13:15
JFK	LAX	B763	10:30	13:23
JFK	MCO	B752	10:25	13:11
JFK	PAP	A306	10:00	13:53
JFK	PHX	A320	10:00	13:24

Frequency (step 4) **Block Time (step 4)**

Timetable *input* strongly influences the resulting modeled *output*. For example, modeling an airport with only ten scheduled operations a day will produce drastically different results than when modeling the same airport with 1000 scheduled operations a day. The timetable determines the level, schedule intensity (i.e., peaky vs. distributed) and general directional flow of these operations. Using a realistic timetable of aircraft operations is therefore a critical component to the operational modeling effort.

In the past, timetables have been derived by extrapolating counts and forecasts of airport terminal operations into individual flights. The FAA measures overall NAS traffic in terms of annual operational counts at each terminal. It publishes a forecast of this traffic each year in the Terminal Area Forecast (TAF). The TAF makes a logical data set to use for growing any hypothetical schedules as well. The allocation of terminal area arrivals and departures to specific airport pairs is accomplished by using what is known as a Fratar algorithm [3]. The downside to this method, however, is that in

[1]The NAS is a large network of airports and air traffic control facilities (ATC). ATC facilities are classified into three categories: airport towers, terminal radar approach control facilities (or, TRACONs), and air route traffic control centers (ARTCCs or, en route centers). Towers are located at airports and direct airport traffic on the ground and within approximately 5 nautical miles of the airport to altitudes of about 3000 feet. There are 496 towers, of which 266 are under FAA direct control and 230 are managed under contract. TRACON facilities sequence and separate aircraft as they approach and leave airports beginning approximately 5 nautical miles and ending approximately 50 nautical miles from the airport and at altitudes up to about 10,000 feet. En route centers control aircraft in transit and during approaches to TRACONs. The airspace that most en route centers control extends above 18,000 feet for commercial aircraft. At present, there are 22 en route centers.

reality NAS operations are made up of flows which vary with geography and time. This is not necessarily equivalent to extrapolating terminal area growth into operations. Thus, the challenge is to move from a generic traffic count at a specific terminal to a timetable of flights that includes a "when" and a "where" dimension.

The top-down approach underlying the allocation described above achieves its goal of replicating the intended volume of flights at each airport (i.e., predicted TAF levels), but it does not necessarily achieve the desired operational level of integrity. In other words, the existing method is not capable of forecasting route-specific growth in operational flows. By simply matching terminal traffic forecasts with the current airline schedules found in the Official Airline Guide (OAG), the existing method clearly misses out on rich geographic growth patterns.

The information that is missing in this process is the origin and destination of the passengers on the flights, and information about the routes over which they fly. For example, if some city pairs are expected to experience above- or below-average growth, then some routes (and thus some specific airport pairs) will experience above- or below-average growth. A similar story can be told for hub airports, some of which may add capacity in the near future, and others of which may remain capacity constrained.

CONCEPTUAL FRAMEWORK

At CAASD, we have built a framework which attempts to fill the gaps mentioned above using a bottom-up, origin and destination (O&D) demand-driven microeconomic approach [4]. Our ultimate goal is to produce a timetable of flights that is linked with O&D operations via passenger route choice and carrier equipment choice. (Table 1). Our output should thus be consistent with the OAG, but not driven by it. Our method is comprised of six basic steps.

The first step of the process lays the foundation upon which all the other steps will be built. This step estimates the economic and demographic drivers for where people ultimately want to travel. This is done by econometrically estimating the demand for people's travel between O&D Metropolitan Statistical Areas (MSAs). Passenger O&D trips are specified as a function of average fare, local area income, population, and various market characteristics (Figure 1). Passenger trips are then forecast by combining our model with commercially available forecasts of metro area population and income.

Once we have a model of O&D passenger demand, we use a discrete choice model to determine what itineraries they will choose to get them there. This second step produces the flight segments that will be listed in our timetable. For instance, a person planning to travel from Seattle to New York may have a stopover in Chicago. This process translates a single trip into two separate flights, one from Seattle to Chicago, and the other from Chicago to New York.

The third step determines what type of aircraft will be flown on each flight segment. For instance, the flight from Seattle to Chicago may require a different type of plane, because the distance is twice as great as the distance from Chicago to New York.

Step 4 assigns arrival and departure times to these scheduled commercial flights. Finally, steps 5 and 6 account for flight activity that is not driven by domestic scheduled passenger demand; this includes cargo, international, and general aviation (GA) flights. While the scheduling methodology does not encompass all of the complexities airlines must account for when creating their actual schedules, it does attempt to capture the same passenger demand element that is the primary driver.

METHODOLOGY

1. Estimating and Forecasting Domestic O&D Passenger Demand

People fly because they want to go to places for business and leisure. These decisions are primarily driven by local economic and demographic characteristics. In addition, characteristics such as fare, market share of major carriers, presence of low-cost carriers, seasonality, and the structure of airport hubs all play important roles in eventually determining the O&D demand. Differentiating the NAS by distances, we estimate a set of econometric relationships that define these relationships on O&D data [4]. Figure 1 describes the underlying relationships qualitatively:

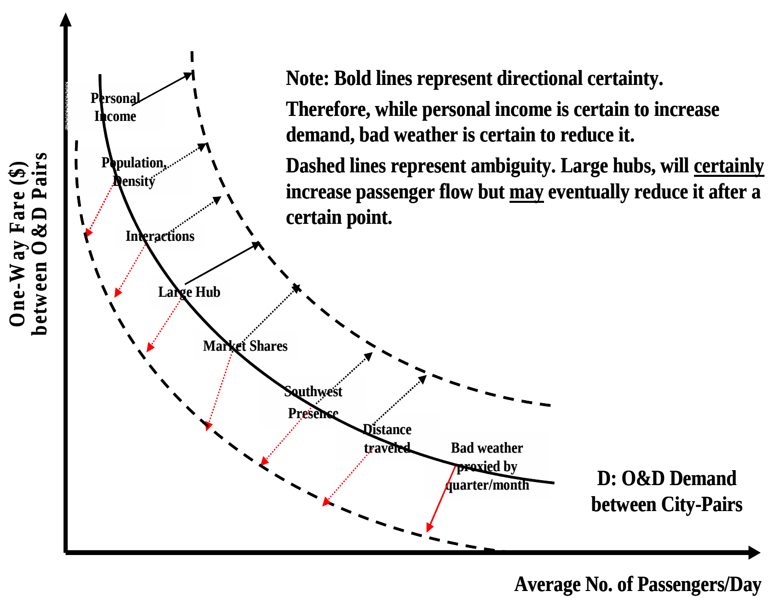

Figure 1. Determinants of O&D Air Demand

To estimate those relationships, data from the Department of Transportation's (DOT's) Origin Destination Survey (i.e., 10% ticket sample) are matched with local economic and demographic information for each O&D airport.[2] Using this information and drawing on well-established econometric methodologies in demand estimation, we estimate O&D passenger demand between origin and destination city pairs as follows:

$$\begin{aligned}\ln(P_{ij}) = {} & \alpha + \beta * \ln(f_{ij}) + \chi_i * \ln(PI_i) + \chi_j * \ln(PI_j) \\ & + \delta_i * \ln(\text{Density}_i) + \delta_j * \ln(\text{Density}_j) \\ & + \eta * \ln(\text{Market Power}^{D}{}_{ij}) \\ & + \iota * \ln(\text{MarketPower}^{ND}{}_{ij}) + \kappa^{D} * (LCC_{ij}) \\ & + \kappa^{ND} * (LCC_{ij}) + \gamma_i * (\text{hub statusOrigin}) \\ & + \gamma_j * (\text{hub statusDestination}) \\ & + \varphi * \ln(\text{Distance}_{ij}) + \rho * (\text{season}) + \varepsilon_{ij}\end{aligned}$$

where i = origin (O) and j = destination (D); P = O&D passenger flow; PI = personal income of O&D in (1996) real chained dollars; Density = population density at O&D cities. Market power is measured by using a concentration index for dominant carriers (D) and non-dominant carriers (ND) and for low-cost carriers (LCC) separate from the network carriers; hubstatusOrigin and hubstatusDestination are two dummy variables accounting for large hubs (dummy = 1) and not (dummy = 0). Distance is the distance between O and D cities. Finally, errors have been assumed to be distributed normally with mean zero and a constant variance. The symbol "ln" signifies the natural logarithm. [See (4) for detailed discussion on econometric methodology, estimation, and results from earlier estimations.]

Using quarterly data starting from 1995, O&D demand is estimated, on average for 38,000-42,000 markets, based on local metropolitan variables as opposed to national economic and demographic conditions, and hence is called bottom-up demand. Notice that this methodology focuses on estimating O&D passenger flows, as opposed to airport-centric activity, as found in the Terminal Area Forecast [8]. Note also that this is an ongoing process; as new data become available, we re-estimate the above relationships and make appropriate changes in selecting exogenous variables and model specifications.

Finally, by combining these estimated relationships with commercially available forecasts of local economic and demographic variables (provided by Global Insight), we come up with forecasts of passenger flows by O&D metropolitan areas. As evident, when these forecasts are rolled up for a particular airport, they can provide us with activity measures that are comparable to the TAF. [See *Results* section for more details].

2. Assigning O&D Passengers to Routes

We now have forecasts of O&D passenger demand between metropolitan areas. But choosing to travel somewhere is not the only decision a passenger must make. They must also choose how and when they will fly. Unfortunately, data on passenger flows by *day* or *time* do not exist in the public domain. However, the 10% ticket sample does have data on passenger itineraries. This is important, because over one-third of itineraries involve at least one connecting flight. Knowing that a certain number of people want to go from Seattle to New York is only part of the story (as illustrated in Figure 2), and obscures the fact that many of these passengers will change planes in a hub such as Chicago O'Hare or Dallas-Fort Worth. Furthermore, flights through these hubs are filled with passengers going to and from a variety of O&D pairs.

We use the following process to convert O&D passenger flows into airport-to-airport (or "segment") passenger flows. As suggested by [6], each possible route between any given origin and destination is constructed as a set of links and nodes (flight segments and airports). Passengers are then assigned to routes, and at the end of the process, total passenger traffic is summed up for each "link."

There are various ways to assign passengers to routes. Currently, using itinerary data from the 10% ticket sample, we assign passengers based on the most recent historical distribution. Thus, the routes available for passengers to choose from are those which are observed in the actual data. While this precludes the model from "thinking outside the box" to determine other potentially feasible routes, the converse is also true—we do not end up with connecting flights going through small, non-hub airports.

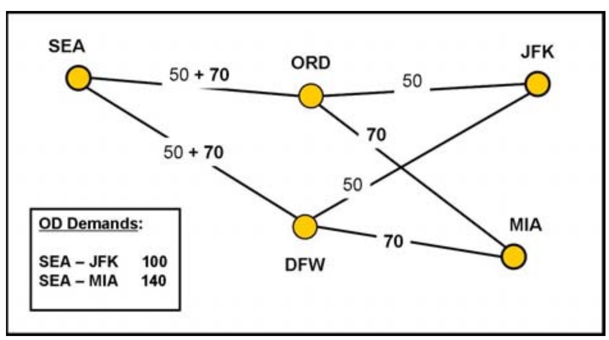

Figure 2. Example of Adding Segment Traffic

A richer assignment method would be to use a multinomial Logit model based on route characteristics

---

[2]Primary data for this analysis is based on the 10 % O&D sample obtained from the Bureau of Transportation Statistics (BTS) [see http://www.bts.gov/oai for details]. In addition, we use T-100 schedule data collected by BTS. We combine the O&D travel data with local economic, demographic and spatial variables collected by the Bureau of Economic Analysis (BEA) (see [4] for more details).

such as travel time, fare, and number of connections. We have experimented with using this type of model, but have not yet incorporated it into our overall system due to the fact that, within our current framework, most of our right-hand-side variables are either static (such as “number of connections”), or assumed static (such as “fare” and “travel time”).

3. Determining Aircraft Equipment Mix

Our third step is critical—taking the airport-to-airport passenger flows and translating total seat demand into the likely set of aircraft types that will fly each route. This step is necessary due to the variation in actual aircraft sizes, which implies that a given number of passengers does not uniquely determine the number of aircraft operations. Estimated passenger counts must be therefore combined with estimated aircraft size to determine a likely equipment mix for each given route.

Choice of aircraft thus emerges as a function of passengers, frequency, trip distance, and other route characteristics. We can therefore estimate a multinomial Logit model that enables us to determine the most likely choice of aircraft type. To do this, we turn to DOT’s T-100 “Segment” data, which combines historical passenger counts with equipment type, along with flight characteristics such as distance. We have examined actual data and classified the vast majority of aircraft based on these characteristics. Of all 311 specific equipment types observed in the second quarter of 2002, there are six natural groupings of aircraft, or categories. Each category has particular performance and capacity characteristics. (See Appendix A and Figure 3).

Category 1 primarily consists of Cessnas and Pipers while Category 2 represents turboprops. On average, these two typically fly segments that are less than 250 miles and, generally speaking, are used to haul passengers between hub airports and small spoke airports. Category 3 consists of regional jets (e.g., ERJs and CRJs) that fly an average distance of 250-500 miles.

Short-haul narrow-bodies (Category 4) generally have a seat range of 80–162 passengers, best cruise at speeds of 550–625 miles per hour, and are observed to have been in maximum use for flights ranging from 500 to 750 miles. Long-haul narrow-bodies (Category 5), on the other hand, have a seat range of 130-190 passengers, best cruise at speeds of 600+ miles per hour, and mostly are in use for flights ranging from 750-1500 miles. The 6th category is the wide-body category that flies the longer haul flights (e.g., 747, 767, 777, L-1011, A300, A310, etc.).

As noted, these distances are averages, and many aircraft within one category also travel distances defined under other categories. Each category is also associated with an average number of seats (passenger capacity) and a best cruise speed. Together, the six classifications account for almost all the scheduled passenger activities (see Figure 3). The most utilized aircraft in the NAS have been narrow-bodies (Categories 4 and 5), which transport approximately 80% of scheduled passengers.

Modeling aircraft choice based on historical passenger activities is a tricky task, especially during a transitional time, such as the one resulting from the terrorist attacks on September 11, 2001 and the economic slowdown preceding 9/11. The U.S. aviation industry is currently undergoing serious structural changes. At the end of 2002, more than 1400 aircraft were temporarily parked in the Mojave Desert.

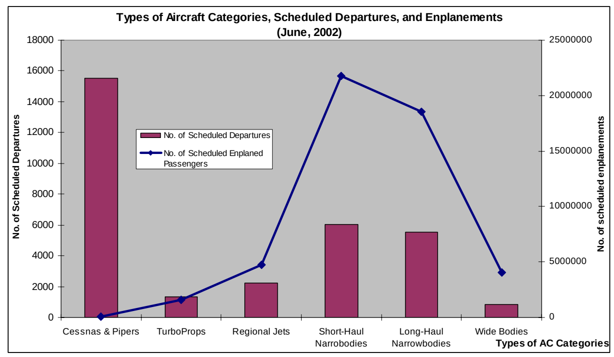

Figure 3. Aircraft Categorization

This is a relatively large percentage considering that the current aircraft inventory consists of only 3623 wide- and narrow-bodies, and around 1020 regional jets (RJs). Routes are being rationalized based on individual profitability in an attempt to improve aircraft utilization. Despite all these changes, the central element of scheduling flights remains intact: carry passengers between two points using the most efficient aircraft given all other characteristics. Thus, we postulate the following multi-nomial Logit equation in order to capture the likelihood of aircraft choices underlying any segment [for more details, see 7]:

$$
\begin{aligned}
P_i\,(y_i = j|\, x_i, \beta) \quad &= \quad \alpha_{ij} + \beta_1 * (\text{passengers}) + \\
(j = 1, 2, \ldots, 6) \quad & \qquad \beta_2 * (\text{distance}) + \\
& \qquad \beta_3 * (\text{OriginHubDummy}) + \\
& \qquad \beta_4 * (\text{DestinationHubDummy}) + \varepsilon_i
\end{aligned}
$$

where j=1, 2, …. 6 represent six different aircraft category choices; passengers are the segment-pair passengers; distance is the distance between two segment points; and OriginHubDummy and DestinationHubDummy are dummy variables representing the airports if they were large airports from where the segment flight originated or landed. We use maximum likelihood (ML) estimation procedure for estimating the above multi-nomial Logit choice model.

4. Assigning Times for Scheduled Domestic Flights

Our fourth step determines when the flights will occur. The first part of this task is to determine exactly how many flights will occur. To do this we combine passenger movements between airport pairs (determined in step 2), with the aircraft that are predicted to fly between those airports (determined in step 3), and then apply a load factor. The load factor applied is derived from BTS data. For example, if 1000 passengers are predicted to fly between LaGuardia and O'Hare on a given day, in a category 4 aircraft (that holds approximately 133 passengers) with a load factor close to 75%, the resulting frequency is 10 flights a day. This task is critical to the calibration of our model since the frequency calculation determines the total level of operations at an airport. By fine tuning the frequency calculation during the model validation stage, the overall number of operations in the timetable can be adjusted up or down to make it more accurate.

The next task is to take the actual flights and assign arrival and departure times. The timetable will contain commercial operations for 292 airports. The number of airports in the timetable with unscheduled or general aviation (GA) activity will be considerably larger (see step 6). Both commercial and unscheduled departure and arrival times are then assigned using historical data, when available.

Historical data for commercial traffic is obtained from the OAG. To get arrival and departure distributions, we use OAG data from 5 different historical years: 1995, 1997, 1999, 2001 and 2002, and the current year of 2003. These data are then transformed into arrival and departure distributions of operations over a given day. Several different "days" were obtained; one weekday, and one weekend, from each of the four quarters, for a total of eight representative days. These different "days" represent specific patterns in seasonal and weekday passenger travel.

Current baseline OAG operations are used as the timetable starting point for scheduled times between city pairs. These data are processed and altered to accommodate changes in forecasted equipment and international traffic. The historical airport operational distributions are used to determine time assignment for additional flights. For example, if eight flights currently operate between LaGuardia and O'Hare each day, and ten frequencies are predicted, then the two additional flights will need to be assigned departure and arrival times based on historical data. The initial eight flights receive the departure and arrival times that already exist in the current OAG schedule, with some minor adjustments.

To assist in the process of assigning departure and arrival times, the airports have been grouped into four tiers. The airports with the most dominant schedules were assigned to tier 1 (the FAA's capacity critical airports), and the airports with the least dominant schedules were assigned to tier 4 (typically airports only served a couple times a day by just one carrier), with the other airports falling somewhere in between. This categorization helps us to determine flight's scheduled arrival and departure times .

Once an arrival or departure time is established at an airport for a given flight, there is very little "slack" left in the schedule on the other end because the equation—departure time, plus block time[3], equals arrival time—is fairly tight. This causes a problem when a departure is created at an airport that then dictates an arrival at the destination airport during an unlikely time or vice versa.[4] To mitigate this problem, an airport's tier is used to determine how strictly additional arrivals and departures should conform to the airport's historical distribution of flights. For instance, a hub airport traditionally has strong arrival and departure banks, thus additional flights should conform to those times when banks occur. On the other hand, adding flights to small spoke airports that have only a handful of operations should not necessarily comply with a historic distribution (i.e., if there are only five flights a day at an airport and a sixth is added, that sixth flight should not necessarily be added at a time when flights have historically occurred). Hence when scheduling between tiers, the more dominant airport's schedule prevails. Intra-tier flights (i.e. flights to and from tier one airports) undergoes additional logic in order to find a departure and arrival time that fits with the airport's historical distributions. This preference is also enforced by the order in which these flights are assigned in the timetable. Flights between tier 1 airports are scheduled first, followed by flights between tier 1 and the remaining tiers. The flights between tier 2 airports are scheduled next, and so on, until all commercial flights have been scheduled.

5. Adding in Scheduled International Flights

Step five adjusts our tentative timetable by taking into account aggregate flows of international passenger and cargo traffic to and from the continental United States (CONUS). Although our modeling focus is the CONUS, we must still account for the additional terminal area traffic, especially important for the international gateway airports, that is generated by flights with only one of their

[3] Block time, the time it takes to leave the gate at one airport and arrival at the gate of the destination airport, is largely a function of distance, winds and aircraft speed, but can also be highly influenced by taxi-in and taxi-out times at the arrival and departure airports respectively.

[4] Some time frames are also unavailable due to constraints on airport operating hours.

endpoints in the CONUS. This is accomplished using a modified top-down approach. All non-CONUS destinations are associated with some pre-determined growth rates (e.g., FAA, IATA, etc.). These rates are then applied to the number of seats currently being flown to or from those destinations. Applying the growth rate to the number of seats is in line with a passenger demand focus, and also allows for smaller increments of growth.

In 2000, around 26 million passengers traveled to the U.S. from around the world. While a majority (43%) of these passengers originated in Western Europe, the Far East had a respectable 29% share, followed by South America's share of 11%. Almost all of this traffic takes place through 11 gateway airports in the U.S., and therefore greatly influences the schedule at these airports.

Unlike our O&D model for domestic air travel, here we use regional growth rates from external agencies (e.g., U.S. Department of Commerce International Trade Agency (ITA), FAA, IATA) to drive our forecasts of international passenger travel. Using these forecasts and assuming the types of aircraft that are currently flown between these destinations, we derive the forecast demand for scheduled departures and arrivals. These data are then added to our domestic schedule.

6. Adding in Non-Scheduled Flights

The last step is to account for unscheduled, or GA traffic. Both the terminal and TRACON handle a large amount of GA traffic. There were an estimated 218,000 active GA aircraft in the NAS, which flew almost 40 million operations in 2000 [8]. Almost four-fifths of this traffic was in the domain of VFR[5] and thus less likely to crowd the en route air space. However VFR traffic impacts airport towers and TRACONs the same as IFR traffic. Given its significant utilization of NAS infrastructure, we must include a model of GA traffic in our timetable.[6]

GA traffic is comprised of many different types of operators. Unscheduled business operators tend to file and fly IFR flight plans. As in the case of scheduled O&D travel, particularly the premium fare travel, unscheduled IFR flights can be sensitive to economic or financial factors. Specific location and time data for these flights are derived using the FAA's Enhanced Traffic Management System (ETMS) data. A breakdown of our IFR traffic forecast is shown in Figure 4.

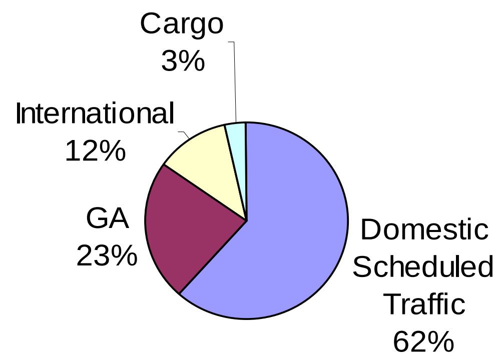

Figure 4. Distribution of Forecast IFR Traffic by Type

On the other hand, data on GA traffic that fly VFR flight plans, as well as military traffic, are not time and location specific, and thus individual operations must be derived using the top-down approach described in the introduction. Based on the historical trends and economic factors, composite growth rates are applied to both VFR and IFR operations to produce forecasts of activity.

RESULTS

Primary results arising from this analytical framework, as noted earlier, are future timetables driven by O&D traffic flow forecasts (i.e., passenger enplanements and aircraft operations). In many ways, these future timetables resemble the Official Airline Guide. One difference, however, is that our timetables also include entries for unscheduled flights, both IFR and VFR, while of course the OAG does not. Nonetheless, comparisons can be made between the OAG and the scheduled component of our timetables.

We identify three areas for comparison with the OAG: first, the distribution of arrivals and departures at an airport at a particular time; second, segment traffic between airports implied by the route allocations; and third, the count of daily flights between airports, i.e., rolled up terminal activities. In order to make these comparisons, we undertook a one-year forecast using estimates based on historical data (1995:Q1 – 2002:Q2). Using the OAG from the second quarter (April), 2002 as our baseline, we forecast schedules for the second quarter, 2003. This forecasted schedule is then compared against the actual OAG for April, 2003.

---

[5] VFR stands for Visual Flight Rules. All scheduled commercial aircraft are required to fly Instrument Flight Rules (IFR). GA or unscheduled traffic can fly VFR or IFR,

[6] Notice that a GA timetable is somewhat *fictitious*. GA traffic, for all practical operational purposes, is unscheduled traffic and hence does not produce timetable on its intent to fly between cities at a given time. It is worth noting here that published schedules are different than the flight plans that IFR GA flights are required to submit. Creation of a schedule, based on their behavior modeled using economic and/or other logic, runs the risk of being truly unreal. Nonetheless, we proceed with this method because of our need to model this entity in simulations of the NAS where they compete with scheduled commercial and non-commercial traffic for scarce air space resources.

1. Arrival and Departure Times by Airport

We constructed an index to validate and verify our results. Notice that there are two ways that our forecasts can deviate from the OAG actual: first, the distribution of arrival and departure banks may differ causing deviation; and second, the number of aircraft operations implied by those banks may differ from that of implied by the OAG. In order for our schedules to be perfectly aligned with that of observed/actual OAG, we require that both distributions of aircraft operations and arrivals and departures will have to be correct simultaneously. This is a rather stringent measure because getting either one of the two is relatively easier than getting both right simultaneously. Nonetheless, we report the cumulative index measuring all deviations.

Generally speaking, total deviations for the top 34 CONUS (i.e., contiguous US) airports were fairly low. While top 34 airports accounted for more than 63% of all operations in the NAS, they accounted for only 4% of the hourly deviations. In other words, most of our deviations occurred at smaller airports. In aggregate, 87% of all forecasted operations came within 25% of actual results.

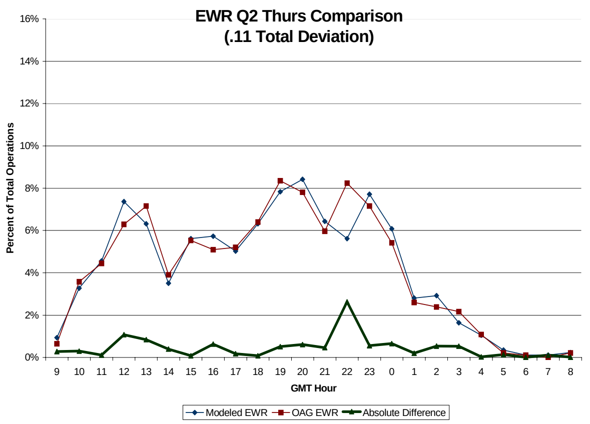

Figure 5. OAG actual schedule and forecast schedule for Newark International (EWR) Airport

As Figure 5 demonstrates, the absolute value of deviations between the actual OAG and the forecast schedules for Newark (EWR) turns out to be fairly low, amounting to a cumulative total of .11. A cumulative index value = 0 implies that there is no absolute difference between the two schedules, i.e., arrival and departure distributions along with number of operations perfectly matched the actual OAG while a value = 2 would imply complete deviations, i.e., time distributions and/or distributions of number of aircraft operations match the actual.

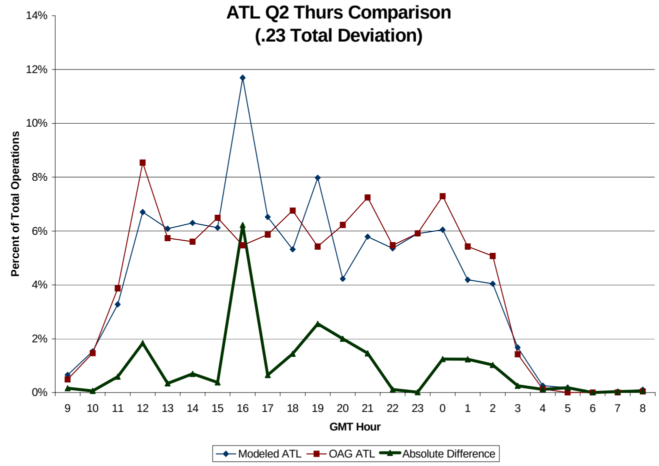

Figure 6. OAG actual schedule and forecast schedule for Atlanta's Hartsfield (ATL) International Airport

Atlanta's Hartsfield International Airport (ATL), on the other hand, scores a relatively higher deviation, 0.23. The highest deviation scores of 0.46 have been recorded for the Washington Dulles International Airport (IAD) and Fort Lauderdale Hollywood International Airport (FLL) while lowest deviation score (0.10) was recorded for Phoenix SkyHarbor International Airport (PHX).

2. Daily Flight Counts by Airport Pair

A second comparison that can be made is to compare our 2003 forecasts of daily flight counts by airport pair relative to the same counts from the actual OAG. Of course, there are several thousand airport pair combinations, but we can summarize our results by looking at the distribution of the absolute value of the differences between the forecast and the OAG. This is shown in Figure 7. As can be seen, there are a few instances of relatively high errors, but the vast majority of our forecasts deviate from the actual OAG by only 1 or 2 flights per day.

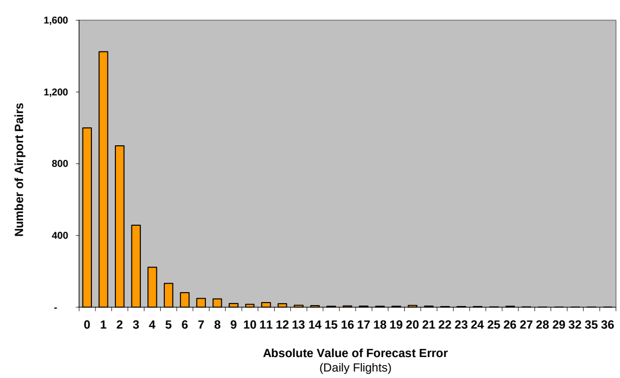

Figure 7. Absolute deviations of forecast flight counts versus OAG flight counts

3. Comparing Terminal Counts with the TAF

As noted earlier, if we sum up our O&D forecasts by airports, we can also arrive at terminal activities, both passenger enplanements and aircraft operations. These can then be compared with the FAA's Terminal Area Forecast (TAF) (8).

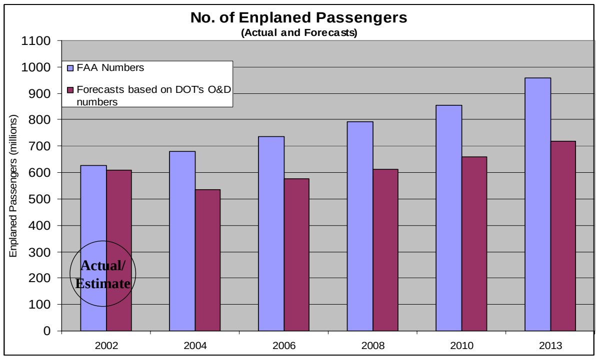

Figure 8. Forecasts of Enplanements: FAA and O&D

As Figures 8 and 9 demonstrate, forecasts based on the O&D methodology described above and those derived from the FAA's TAF compare fairly well. As also can be seen, the O&D forecasts of enplanements and aircraft operations are consistently lower than those of the FAA. However, some of this is due to the fact that our O&D model currently covers only 292 of the airports in the NAS.

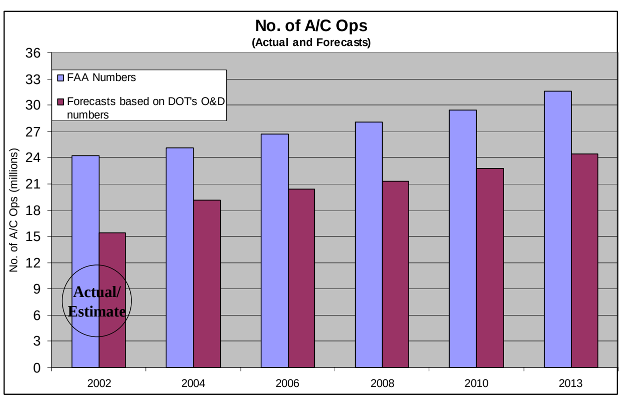

Figure 9. Forecasts of Aircraft Operations: FAA and O&D

CONCLUSION

This paper lays out both a conceptual and an empirical framework to estimate and predict future aviation timetables. Starting from O&D flows of passengers that are determined by local area economics and demographics, existing route allocations are used to derive aviation activities at the terminal levels. Using forecasted values of economic and demographic variables for local metro areas, we derive forecasts for O&D passenger flows. Passenger itineraries are then assigned using the existing allocations as a guide for predicting the future network. Aircraft type and the number of operations between airports have been derived by using a multinomial aircraft choice model that ties enplanements to aircraft choices. Once flight counts have been forecast, we calibrate the model to align with the baseline OAG. After this calibration has been established, we use the forecasted values of O&D flows to derive the future aviation timetables. The final steps then use a top-down method to estimate international, cargo, and unscheduled traffic counts and times.

The primary goal of this work is to support modeling efforts that evaluate NAS performance. As evident, the present framework makes significant qualitative improvement over existing work in that comparative static analysis can be done fairly easily under our framework. That is, by changing the inputs that go into the process, our framework can be used to perform various types of "what if" policy analysis, and thus can stand on its own as a useful analytical tool. For example, changes in demand conditions, route allocations and airport share and their impact on schedules can be incorporated fairly easily in our framework. Distribution of weights on time schedules, relatively more weight on current schedules (i.e., 2003) as opposed to history (i.e., 1995-2002), can be changed and the impact studied in our framework. Similarly, weights on airports can be changed in order to understand the impact on schedules. In our current framework, larger airports receive higher weights in determining both arrival and departure banks. However, this may change in the future. All these parameters can be adjusted to generate alternate schedules for validation and/or scenario analyses.

REFERENCES

1. Wieland, F. 1999. "The Detailed Policy Assessment Tool (DPAT)," MITRE Technical Report MTR99W00000012, McLean, VA: The MITRE Corporation.

2. Wieland, F., C. Wanke, B. Niedringhaus, and L. Wojcik. 2002. "Modeling the NAS: A Grand Challenge for the Simulation Community." Proceedings of the Grand Challenges in Modeling and Simulation Conference, San Antonio, TX: 2002 SCS Western Multi-conference.

3. Long, D., D. Lee, E. Gaier, J. Johnson, and P. Kostiuk. 1999. "A Method for Forecasting Commercial Air Traffic Schedule in the Future," NASA/CR-1999-208987, Hampton, VA: NASA/LMI, Langley Research Center.

4. Bhadra, D. 2003. "Demand for Air Travel in the United States: Bottom-Up Econometric Estimation and Implications for Forecasts by O&D Pairs,"

Journal of Air Transportation, Vol. 8, No. 2(2003) pp. 19-56; paper also presented at the 6th Air Transport Research Society Conference (ATRS 2002) hosted by the Boeing Corporation, July 14 – 16, 2002, Seattle, Wa; and AIAA/ATIO Technical Forum, October 1–3, 2002, Los Angeles, CA.

5. General Accounting Office. 2002. "Commercial Aviation: Air Service Trends of Small Communities Since October 2000, Report to Congressional Reporters," GAO Report No. GAO-02-432, Washington, DC.

6. Oppenheim, N. 1995. *Urban Travel Demand Modeling: From Individual Choices to General Equilibrium*. New York: John Wiley and Sons.

7. Bhadra, D. 2003. "Choice of Aircraft Fleets in the US NAS: Findings from a Multinomial Logit Analysis", Paper to be presented at the 3rd ATIO/AIAA conference, Nov. 17-19, 2003, Denver, Colorado.

8. Federal Aviation Administration. 2003. "Aerospace Forecasts: Fiscal Years 2003-2014," Washington, DC: U.S. Department of Transportation.

Paper presented at the 13th Annual Federal Forecasters Conference at the BLS Training Center, Oct. 27th, 2003. An earlier version of this paper was presented at the NEXTOR-FAA *Conference on Air Traffic Management and Control,* Virginia Tech Graduate Center, Falls Church, Virginia, June 2-3, 2003; and Transportation Simulation session of the Annual Summer Computer Simulation Conference (SCSC'03), Montreal, Canada, July 20-24, 2003.

Appendix A:

AIRCRAFT CATEGORIZATION AND RELATED INFORMATION

	Types of Aircraft	Representative Types of Equipment in this Category	Av. Distance and Ranges	Avg. Size Range Pax (min)	Pax (max)
Category 1	**CESSNAS AND PIPERS**	CES-150 - 185; AERO-200; PA 12-32	150	3	20
		Overall numbers with upper and lower bounds	150	3	20
Category 2	**TURBO PROPS**	ATR-42 Aerospatial; ATR-72 Aerospatial; Dornier 328 Turbo;	< 250	30	37
		TurboProp 1-2 engine; JETST-31 BAE; JETST-41 BAE;	< 250	60	72
		Overall numbers with upper and lower bounds	250	45	55
Category 3	**REGIONAL JETS (RJs)**	Canadair RJ-100/R;Canadair RJ145-200;	250-500	45	70
		Embraer EMB-135; Embraer EMB-145; EMB-140	250-500	45	70
		Overall numbers with upper and lower bounds	500	45	70
Cateogry 4	**SHORT-HAUL NARROW BODIES**	Boeing B-737-500; Boeing B-737-400;	500-750	127	155
		Boeing B-737-300; Boeing B-737-100;	500-750	105	129
		Boeing B-737-200C; Douglas DC-9-10;	500-750	62	76
		Douglas DC-9-30; Douglas DC-9-40;	500-750	87	107
		MD-80 & DC-9-80; MD-90-30/50;	500-750	135	163
		Douglas DC-9-50; Boeing B-727-100;	500-750	113	139
		Overall numbers with upper and lower bounds	750	105	128
Category 5	**LONG-HAUL NARROW BODIES**	Boeing B-737-800; Boeing B-757-200;	750-1500	155	189
		Euro Airbus A320; Airbus Industrie A319;	750-1500	131	161
		Boeing B-737/LR Boeing B-737-700/; Boeing B-727-200;	750-1500	132	162
		Overall numbers with upper and lower bounds	1500	139	171
Category 6	**WIDE BODIES**	Douglas DC-10-10; Douglas DC-10-30;	1500-3000	278	340
		Douglas DC-10-40; Boeing B-747-100;	1500-3000	256	312
		Boeing B-747-200; Boeing B-747-400;	1500-3000	321	393
		Boeing B-767-200; Boeing B-767-300;	1500-3000	158	194
		Boeing 777; Lockheed L-1011-1;	1500-3000	239	293
		Lockheed L-1011-50; Douglas MD-11;	1500-3000	299	379
		Euro Airbus A-300; Euro Airbus A310;	1500-3000	205	251
		Overall numbers with upper and lower bounds	3000	251	309