

# Understanding Air Transportation Market Dynamics using a Search Algorithm for Calibrating Travel Demand and Price

Vivek Kumar, Brant M. Horio, Anthony H. DeCicco, Shahab Hasan, and Virginia L. Stouffer  
*LMI, Tysons, VA, 22102*  
*and*

Jeremy C. Smith, and Nelson M. Guerreiro  
*NASA, Langley Research Center, Hampton, VA, 23681*

**This paper presents a search algorithm based framework to calibrate origin-destination (O-D) market specific airline ticket demands and prices for the Air Transportation System (ATS). This framework is used for calibrating an agent based model of the air ticket buy-sell process – Airline Evolutionary Simulation (Airline EVOS) -that has fidelity of detail that accounts for airline and consumer behaviors and the interdependencies they share between themselves and the NAS. More specifically, this algorithm simultaneously calibrates demand and airfares for each O-D market, to within specified threshold of a pre-specified target value. The proposed algorithm is illustrated with market data targets provided by the Transportation System Analysis Model (TSAM) and Airline Origin and Destination Survey (DB1B). Although we specify these models and datasources for this calibration exercise, the methods described in this paper are applicable to calibrating any low-level model of the ATS to some other demand forecast model-based data. We argue that using a calibration algorithm such as the one we present here to synchronize ATS models with specialized forecast demand models, is a powerful tool for establishing credible baseline conditions in experiments analyzing the effects of proposed policy changes to the ATS.**

## Nomenclature

$\alpha$	=	price elasticity of demand
$\lambda$	=	multiplier for airfare of non-stop itineraries
$A$	=	advertised airfare
$B$	=	basefare
$d$	=	number of days to travel date
$l$	=	load-factor (ratio of sold to total number of available seats)

## I. Introduction

**O**NE of the keys to realize the intended benefits of NextGen - the multi billion dollar FAA initiative to modernize United States air transportation system (ATS) – is the participation of airlines who are key stakeholders in this system. The success of almost all NextGen improvements hinge on the willingness of airlines to invest in these new technologies and procedures. Airlines are profit driven entities and can only justify investments that are expected to yield returns. An airline’s profit is not just dependent on its own revenues and costs, but also influenced by the performance of its competitors.

Therefore, in order to realistically model the costs and benefits of any of the NextGen technologies and procedures, it is imperative to model its effect on the economics of airline ticket buying and selling process. However, the mechanism of pricing tickets is unique and proprietary to individual airlines and hence not available in the public domain. The data on realized ticket sales is business-sensitive as well and also not available in the public domain. However, the underlying principles that drive these individual pricing mechanisms, and hence the sale of ticket are governed by the fundamental principles of economics, e.g. the supply-demand theory. Researchers have tried to bridge this gap in the past by developing models that mimic the airlines and consumer behavior in ticket buy-sell process through patterns extracted by trend analysis of publicly available data (e.g.: FAA’s DB1B ticket sample data).

Airline Evolutionary Simulation (Airline EVOS) is an agent based model that was developed to mimic the ticket buy-sell process with a level of fidelity of detail that accounts for airline and consumer behaviors and the interdependencies they share between themselves and the NAS<sup>1</sup>. Although Airline-EVOS is not the focal point of discussion of this paper, for the sake of completeness it is summarized in Section II. The main thrust of this paper is to enable Airline-EVOS for benefit modeling by performing the essential precursor to any ABM execution, i.e. validation and calibration of the agents (in this case the consumers that buy airline ticket), so that the results obtained are credible and representative of the real world scenarios.

This paper is organized as follows. The following subsection details some background information on Airline-EVOS and motivates the objective of this research. The next section discusses the research approach and the algorithm designed for calibration. The following section discusses the results of the approach and finally the conclusion section highlights the findings and the usefulness of this approach in other contexts.

## II. Background Information on Airline Evolutionary Simulation (Airline EVOS)

As mentioned previously, the ticket pricing mechanism for each airline is a business sensitive information, and this strategy for dynamic ticket pricing is unique for each airline. Airline-EVOS is an agent-based model that simulates the ticket buy-sell process by mimicking airlines and consumer behaviors, through patterns extracted by trend analysis of publicly available data. The intended use of Airline-EVOS is to explore how changes—regulatory or technological—affect the airlines, and how airlines, in turn, react and affect the ATS and its stakeholders.

Airline-EVOS employs an agent-based approach to simulating airline behaviors, ideal for the complex system interactions in the ATS. The agent-based approach naturally accommodates modeling the multiple independent, heterogeneous entities, or agents, within the ATS that interact in specific ways. Agents follow predefined and often simplified rules that are designed to emulate real-world counterparts (airlines, passengers, air traffic control, airports). By representing simple behaviors at the agent-level, emergent dynamics can be observed from a system perspective as the agents interact with each other. Key behaviors that airlines are expected to exhibit in response to policy changes and new technologies were identified. At a high level, only two primary behavior responses are directly available to the airlines: airfare pricing strategies and schedule-based strategies. Accounting for this, airline agents are adaptive, making tactical and strategic changes to their airfares and schedules under the influence of reinforcement learning, to best enable themselves to generate profit and compete in markets. Passenger agents choose to purchase their most preferred ticket based on their objectives and mission, subject to some randomness, used to model non-utility maximizing behaviors and irrationality.

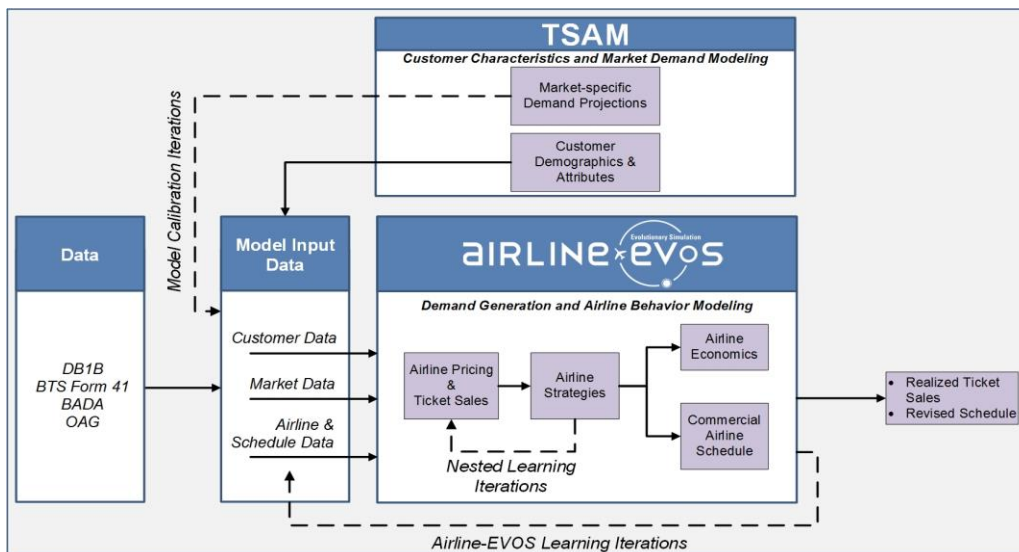


Figure 1: Overview Context Diagram of Airline-EVOS

Figure 1 illustrates, at a high-level, how Airline-EVOS operates, using data inputs for airlines and passengers to execute a looping process within the model. This iterative process models airlines selling tickets and responding to both the market and competition, adjusting airfares or flight schedules to maximize profits. Iteration loops create learning opportunities; the airline agents determine the success of implemented strategies and, on the basis of their

assessment, continue or end the strategy and try new ones. The outputs from Airline-EVOS are adjusted flight schedules, which are used for assessing airline-specific operational performance as well as ATS-wide performance, and airline-specific financial metrics, which are used for assessing airfares, market trends, and airline financial performance.

### III. Data Used

**A. Transportation System Analysis Model (TSAM):** TSAM is a nationwide transportation planning model to forecast intercity travel behavior in the United States. The model uses socio-economic and demographic data to make projections of future travel demand for trips by available modes; these are travel by air, road and rail<sup>2</sup>. TSAM is used for current and future schedule data (2012 and 2030), as well as O-D market specific average airfare and demand. TSAM provides the airfare and demand data as annual aggregate values, which had to be converted to daily level data.

**B. DB1B (Airline Origin and Destination Survey Data) and Form 41 Financial Data (Air Carrier Financial Reports):** DB1B is a 10% sample of airline tickets from reporting carriers collected by the Office of Airline Information of the Bureau of Transportation Statistics (BTS)<sup>3</sup>. Data includes origin, destination and other itinerary details of passengers transported. We use DB1B data to get the (P0,D0) point for each O-D market specific airfare data. The P0D0 point for every market is defined as the number of customers willing to pay at least the 75th percentile of airfare. For this task, the transactions that are non-CONUS and/or have airfares less than \$30 or greater than \$10,000 are removed. These criteria filtered 289,961 of 5,477,362 entries in the original DB1B file used, or approximately 5 percent of all entries.

Form 41 data contains financial information on large certified U.S. air carriers<sup>4</sup>. This financial information includes balance sheet, cash flow, income statement, operating costs etc. We use this data to obtain airline specific financial information to distinguish and account for the unique nature

**C. IATA report (Price Coefficient of Elasticity):** The amount that the individual customers will pay can be determined given the market’s price elasticity curve, derived from own-price elasticity-of-demand estimations based on market haul-length and customer type (business or leisure). Table 1 shows the values of price elasticity of demand ( $\alpha$ ), which are as specified in the report prepared for International Air Transport Association (IATA) by InterVISTAS on estimating air travel demand elasticities<sup>5</sup>.

*Table 1 Own-Price Elasticities of Demand (Source: IATA Report<sup>3</sup>)*

Trip length and traveler type	Low	Mean	High
Long-haul domestic business	-1.428	-1.15	-0.836
Short-haul business	-0.783	-0.7	-0.595
Long-haul domestic leisure	-1.228	-1.104	-0.787
Short-haul leisure	-1.743	-1.52	-1.288

Using airfare and air-carrier cost data from BTS and price elasticities from Table 1 above, multiple instantiations of demand curves are generated. These demand curves are created for each O-D market and travel type (business or leisure) at the beginning of a model run. A sample demand curve for JFK-LAX market and business passengers is shown in Figure 2.

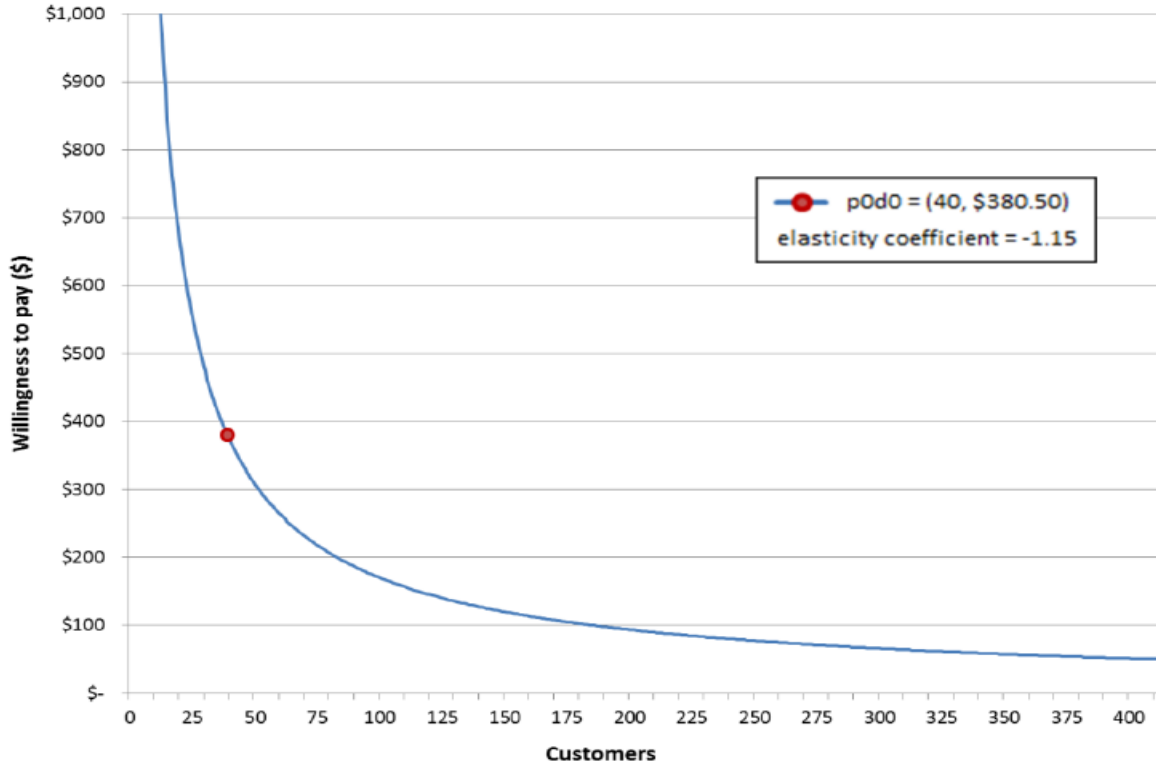


Figure 2 Market Demand Curve for JFK-LAX Using Market Elasticity Assumptions for Business Long-Haul Markets

This demand curve is a graphical representation of the number of potential customers that may buy tickets, as well as the amount each customer is willing to pay for a single ticket. A point (x,y) on the curve implies that there are exactly “x” customers who are willing to buy the tickets if it is priced at “\$y”.

The demand curve is generated using P0D0 values obtained from DB1B and the price coefficient of elasticities obtained from IATA report. Thereafter the following equation was used to iteratively generate the demand curve:

$$P = \left( \frac{D}{D_0} \right)^{1/\alpha} * P_0, \tag{1}$$

P is the willingness to pay at demand D. P<sub>0</sub> and D<sub>0</sub> have been defined earlier in the DB1B data explanation. The value of D is incremented from 0 to D<sub>max</sub> in steps of one. D<sub>max</sub> is arbitrarily chosen to be 1.25 times the TSAM average daily demand for the O-D market under consideration. Consumers are instantiated 25% in excess of target TSAM demand to ensure there are enough customer agents in the ABM ecosystem to satisfy the demand.

### Calibration Objectives and Challenges

The main goal of calibration is to generate representative market-specific demand curves, although it also initializes the model with more representative starting airfares. This process entails shifting the demand curve higher or lower to capture desired demand at the desired average airfare.

One of the major challenges in designing a one-size-fits-all calibration algorithm for all the O-D markets is the interdependence of the airfare and demands for many O-D markets in the ATS. For a given airline, it is possible that multiple O-D markets that share flight legs exist. For example: JFK-ORD-DTW and JFK-ORD-IAD share the JFK-ORD leg. For such market pairs, demand on the common flight leg for one market would influence the remaining supply (and hence airfare and demand) for the other market.

Another challenge in designing a calibration algorithm that also tries to calibrate average airfare for each market is the variation of ticket prices for a given market. This variation in airfares for a given market follows real world scenarios where a particular itinerary for a given O-D market is priced based on number of connections, departure time, and travel duration among other factors. Also, for a given itinerary, the airfares are modified based on dynamic

load factors and advanced purchase time. To account for all these variations in the airfares, Airline-EVOS models airfares (A) for a given O-D market (m) through baseprice (B) and multipliers ( $\lambda$ , F and G) as shown in the following equation.

$${}^i A_m^{l,d} = \lambda_i * B_m * F(l^i) * G(d^i) \quad --(2)$$

${}^i A_m^{l,d}$  is the airfare for an itinerary ‘i’ of market ‘m’ when the load factor on itinerary ‘i’ is  $l^i$ , and  $d^i$  is the days-to-travel for itinerary i. F is a function that determines based on the input load factor, the airfare multiplier<sup>6</sup>. Similarly, G is a function that determines based on the days-to-travel, the airfare multiplier<sup>6</sup>.  $\lambda_i$  is a non-stop multiplier that is 1.5 if itinerary i is non-stop and 1 otherwise.

During calibration, B is the only parameter that can be adjusted to influence airfare. The dynamic multipliers related to load factor and advance purchase time makes it particularly challenging to alter basefare such that the average of resulting airfares is within tolerance of the TSAM target airfare.

Calibration is successful when the total number of customers that travel in Airline-EVOS matches—within some accuracy threshold—the number of customers that travel in TSAM, for each market, and simultaneously, the average price of the two models are within some accuracy threshold. In Airline-EVOS, during the customer ticket purchasing phase, customers purchase tickets in a sequential order based on their advance purchase date. As each customer buys a ticket, airlines change their ticket prices. For the most part, the more tickets sold, the higher the price becomes since price increases with load factor. There can be exceptions to this rule since price also changes with advanced purchase time. Regardless, price is unpredictable and changes with the number of passengers. Therefore, it is impossible to simply modify customer WTP values until the traveler counts match between the two models. Instead, to calibrate the Airline-EVOS demand model to TSAM, it is necessary to employ a search algorithm that will increase and decrease both base price and customer WTP values until the values for average price and number of trips match between TSAM and Airline-EVOS.

#### IV. Method

The search algorithm proposed in this paper adjusts, for each market, the base market airfare and the customer WTP, in order to match the simulated Airline-EVOS demand to match the market-specific TSAM average demand at the TSAM average airfare.

The following is the description of the method:

From the 54,550 distinct origin-destination (O-D) markets modeled in TSAM, a market was filtered out if that particular market satisfied any one of the following conditions:

- i. Has no available itinerary option (as per input schedule derived from TSAM),
- ii. Has an average airfare less than \$30, or
- iii. Has a daily demand projection of fewer than 20 passengers.

For each market remaining in the input data (4,688), it is considered to be calibrated if the average Airline-EVOS demand at the TSAM average airfare is within 30 percent of the TSAM target. The markets in the ATS are highly heterogeneous, both in terms of airfare and demand. The high variability in the range of these target values makes it challenging to design a one-size-fits-all calibration algorithm solution. This challenge is exacerbated by the fact that this is a network problem where these markets are not disjoint, and the behavior of one market influences other connected markets. Considering all these factors, the calibration tolerance limit for both the market demand and the airfare was agreed up with NASA as 30 percent; a market was considered successfully calibrated if its demand and airfare were within 70 percent to 130 percent of the TSAM target value.

The calibration algorithm approach used two phases to address the two challenges described earlier. The first phase was for initial calibration of Airline-EVOS to within the accuracy tolerance of the TSAM target. The second phase was an evolution phase in which airlines used TSAM calibrated parameters for market demand curve assumptions and was then allowed to conduct some number of learning iterations to evolve their own airline-specific base prices. That approach enabled each airline to have its own valuation of a given market, which accounted for factors such as operating costs, market share, and fuel costs. This was necessary to allow differentiation between airlines.

This two-phased approach yielded a TSAM calibrated result that (1) captured a representative sample of market demand dynamics, while (2) also accounting for airline-specific pricing behaviors within that market.

##### A. Phase 1 Initial Calibration to TSAM

The purpose of this phase is for the calibration algorithm to determine some market-specific customer WTP and base market price—for all airlines—that results in an Airline-EVOS average market demand that closely resembles the TSAM average market demand at the TSAM average airfare. It ensures a better initial condition, moving Airline-EVOS parameters closer to the TSAM target values, and thereby helps the overall calibration algorithm to converge more quickly. Our search algorithm for Phase 1 is presented in the following pseudo-code:

1. For each selected market, initialize
  - a. base fare to 50 percent of the target TSAM average airfare for that market, and
  - b. demand curve such that exactly 120 percent of TSAM demand passengers' WTP is less than or equal to the base fare.
2. Initiate the ticket purchasing component of the simulation to compare Airline-EVOS simulated results against the TSAM target values for each market.
3. Terminate the calibration loop if either of the following is true:
  - a. Calibration iteration reaches the user-specified maximum number of iterations. After 6-9 iterations, diminishing returns were observed, where each iteration executes in about 3 minutes.
  - b. Calibration results are within the accuracy target of 30 percent of TSAM target values.
4. For each market that is not calibrated, with respect to demand and/or air-fare, do the following:
  - a. If demand is not calibrated:
    - i. If demand > TSAM demand, then decrease WTP.
    - ii. Otherwise, increase WTP.
  - b. Else (implies that only airfare is not calibrated):
    - i. If average airfare > TSAM airfare, decrease base airfare.
    - ii. Otherwise, increase base airfare.
5. Go to Step 3 and continue calibration.

## **B. Phase 2 Evolution of Airfares**

The search algorithm used in Phase 1 modified a market-specific base airfare used by all airlines operating in that market. The resulting Airline-EVOS outcomes from this phase did not preserve those agent characteristics that distinguished the airlines from each other. Therefore, a secondary calibration step was initiated to allow airlines to evolve their market-specific base fares (keeping the calibrated WTP values static) that accounts for their own costs, revenues, and competitive position. This phase is based on a learning algorithm that seeks to improve the airline's itinerary-based profits in that market.

Given outputs from the first calibration phase, airlines are now allowed to evolve only the base airfares, keeping the market-specific demand curves (WTP) intact. The following pseudo-code describes our second phase:

1. Initiate the ticket purchasing component of the simulation to compare Airline-EVOS simulated results against the TSAM target values for each candidate market identified in Phase 1 of the calibration.
  2. Terminate the calibration loop if learning iteration number is equal to the user-specified maximum number of iterations, currently set at 10, else go to Step 3.
  3. Instruct all airline agents to calculate profits on an itinerary basis for each market.
  4. For each airline in a candidate market:
    - a. If the itinerary-based profits have improved from the previous state, increase the learning factor multiplier and apply it to the base fare, increasing the airfare.
    - b. Otherwise, decrease the learning factor multiplier.
- Go to Step 1 and continue process.

## **V. Results**

This calibration algorithm was applied and tested on 2012 and 2030 flight schedules.

### **A. Current Schedule (2012) Calibration Results**

Phase 1 resulted in the calibration of 3,084 of the 4,688 candidate markets after six iterations. Another three learning iterations increased the total to 4,309 of the 4,688 candidate markets, which represents about 74.3 percent of total TSAM passenger demand. We regarded this as sufficient for the initial starting parameters and representation of market dynamics to begin the experiments.

Figure 3 shows the results over the nine iterations we used for calibration. The top plot shows the calibration to TSAM average demand, and the bottom plot shows the calibration to TSAM average airfare. The y-axis measures calibration error in both plots and reflects the difference between simulated Airline-EVOS results and the TSAM

target values, divided by the TSAM target. We take the absolute value of this ratio, because our results may be on either side of the TSAM target. The smaller this ratio, the more the simulated results match TSAM. We consider a market as calibrated when the ratio for measures of both demand and airfare is less than or equal to 0.30. As typical of box plots, the horizontal red line in the box represents the median, but we have also overlaid an open circle to represent the mean. Table 2 shows the total number of markets calibrated at the end of each iteration.

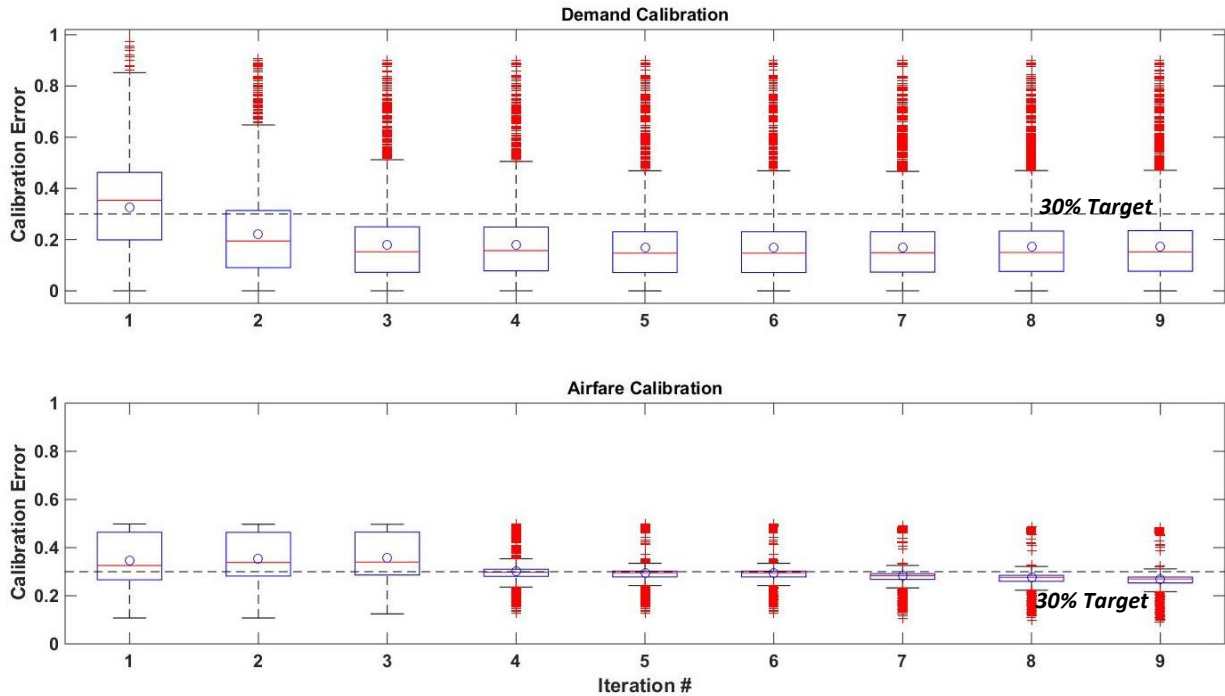


Figure 3 Airline-EVOS Calibration Error Results for 2012 Schedule

Table 2: Airline-EVOS Markets Calibrated Results for 2012 Schedule

Calibration Phase	Initial Calibration Phase						Market Learning Phase		
Iteration	1	2	3	4	5	6	7	8	9
Markets calibrated	594	1,030	1,210	2,412	3,084	3,084	4,265	4,309	4,309

The results shown in Table 2 indicate that after iteration 5, the number of markets calibrated plateaus at 3,084 markets. Market learning is then initiated and improvements in calibration to airfare result in the calibration of 4,309 markets by iteration 9.

### B. Future Schedule (2030) Calibration Results

For the future schedule, a total of 4,905 candidate markets were identified. Phase 1 resulted in the calibration of 3,091 of those markets after six iterations. Another three learning iterations in-creased the total to 4,476 of the 4,905 candidate markets, which represents about 76.2 percent of total TSAM passenger demand. This was considered sufficient to initialize the starting parameters and represent market dynamics to begin the experiments.

Figure 4 shows the results over the nine iterations we used for calibration. The top plot shows the calibration to TSAM average demand, and the bottom plot shows the calibration to TSAM average airfare. The y-axis measures calibration error in the same way as in Figure 3. Table 3 shows the total number of mar-kets calibrated at the end of each iteration.

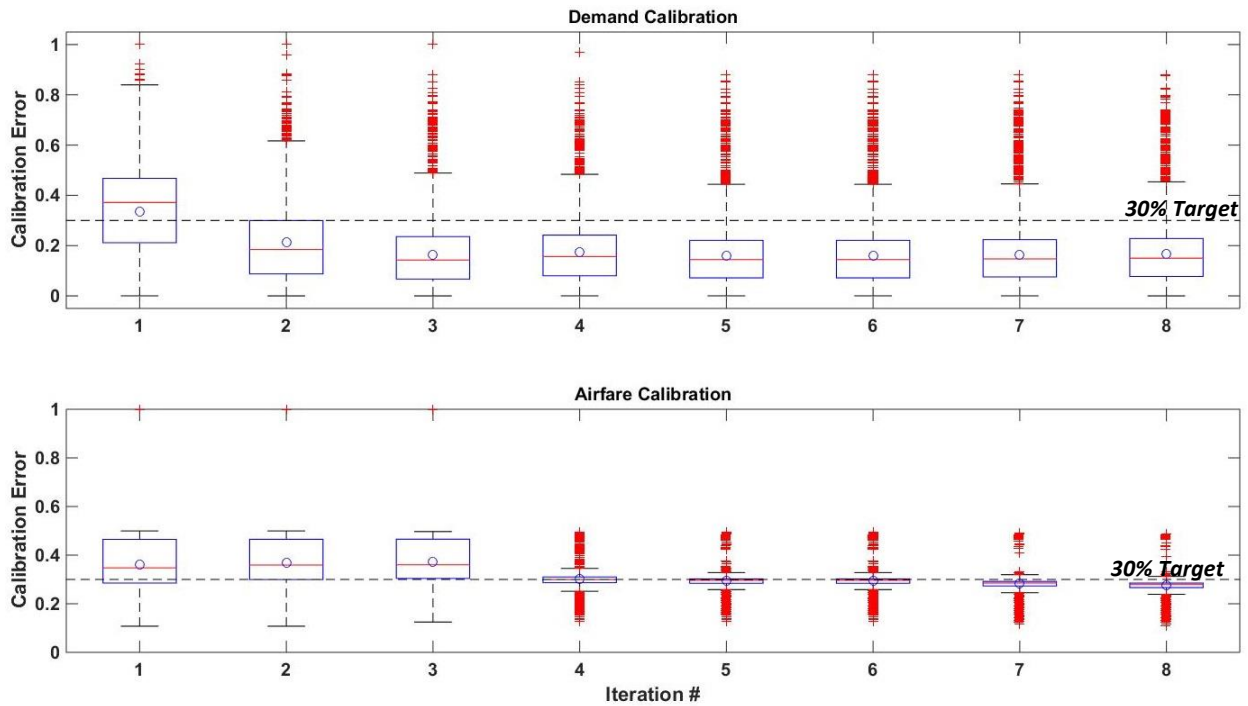


Figure 4: Airline-EVOS Calibrated Results for 2030 Schedule

Table 3: Airline-EVOS Markets Calibrated Results for 2030 Schedule

Calibration Phase	Initial Calibration Phase						Market Learning Phase		
Iteration	1	2	3	4	5	6	7	8	9
Markets calibrated	510	881	1,027	2,329	3,091	3,091	4,420	4,473	4,476

## VI. Conclusions

The search based calibration algorithm demonstrated in this paper has several implications. The main outcome of calibration was to adjust the market-specific demand curves so that the average number of passengers for the given market was within an acceptable tolerance limit of the TSAM-specified demand for that market at a particular average airfare. This calibration ultimately enables:

- (i) Synchronization of the input flight schedule with demand, and
- (ii) Leveraging of TSAM scenarios that model future demand, schedule, and fleet evolution.

The demand curves are sorted by WTP values for each market; individual points on the curve represent potential buyers. By calibrating WTP, we are essentially creating a pool of customers with static WTP values. Each customer is an independent agent in the system, and the customer's decision on whether to buy an airline ticket is driven by (1) a personal utility assessment for price and inconvenience, and (2) the availability of an affordable airfare.

Even though TSAM was used for target demand and airfare values, the proposed algorithm is applicable to calibrating any low-level model of the ATS to some other demand forecast model-based data. Ultimately, it must be stressed that a calibration algorithm such as the one we present here to synchronize ATS models with specialized forecast demand models, is a powerful tool for establishing credible baseline conditions in experiments analyzing the effects of proposed policy changes to the ATS.



## Acknowledgments

This work was accomplished under NASA research announcement (NRA) NNH10ZEA001N, “Research Opportunities In Aeronautics–2010”, Amendment 6, NextGen—Systems Analysis Integration And Evaluation Projected C.3 Subtopic 2: Comprehensive Modeling Of Air Carrier Behavior. The authors would also like to thank William Cotton of Cotton Aviation Enterprises and Vince Costanzo for their counsel and advice, and Jeffrey Goldsmith, James Hebden, and Terry Thompson of LMI, for their contributions and support of this research.

## References

<sup>1</sup>Horio, B. M., DeCicco, A.H., Stouffer, V.L., Hasan, S., Rosenbaum, R.L., and Smith, J.C. “Agent Based Modeling of Air Carrier Behavior for Evaluation of Technology Equipage and Adoption,” *14th AIAA Aviation Technology, Integration, and Operations Conference*. June

<sup>2</sup>Smith, J.C., et al. “Identification and Analysis of National Airspace System Resource Constraints,” *NASA/TM–2015-218695*, March 2015.

<sup>3</sup>Department of Transportation, Bureau of Transportation Statistics, “DB1B Database,” [http://www.transtats.bts.gov/Tables.asp?DB\\_ID=125](http://www.transtats.bts.gov/Tables.asp?DB_ID=125) [cited 16 March 2014].

<sup>4</sup>Department of Transportation, Bureau of Transportation Statistics, “Air Carrier Financial Reports (Form 41 Financial Data),” [http://www.transtats.bts.gov/Tables.asp?DB\\_ID=135](http://www.transtats.bts.gov/Tables.asp?DB_ID=135) [cited 16 March 2014].

<sup>5</sup>InterVISTAS Consulting, Inc., Estimating Air Travel Demand Elasticities: Final Report, prepared for the International Air Transport Association, December 2007.

<sup>6</sup>Hasan, S., Horio, B. M., DeCicco, A. H., Stouffer, V. L., Goldsmith, J. H., Frank, A. B. and et al. Comprehensive Modeling of Air Carrier Behavior: 2013 Annual Report, NS308T3, McLean, VA: LMI, March 2014.