Proceedings of the 2015 Winter Simulation Conference
L. Yilmaz, W. K. V. Chan, I. Moon, T. M. K. Roeder, C. Macal, and M. D. Rossetti, eds.

AN AGENT-BASED APPROACH TO MODELING AIRLINES, CUSTOMERS, AND POLICY IN THE U.S. AIR TRANSPORTATION SYSTEM

Brant M. Horio
Vivek Kumar

LMI
7940 Jones Branch Drive
Tysons, VA 22102, USA

Anthony H. DeCicco

RAND Corporation
1776 Main Street
Santa Monica, CA 90401, USA

ABSTRACT

We present a modeling approach to assist policymakers in identifying impacts on the U.S. air transportation system (ATS) due to the implementation of potential policies and the introduction of new technologies. Our approach simulates the responses of U.S. commercial airlines and other ATS stakeholders to these changes, which cumulatively result in consequences to the ATS. Our research is built upon an agent-based model—called the Airline Evolutionary Simulation (AIRLINE-EVOS)—which models airline tactical decisions about airfare and schedule, and strategic decisions related to fleet assignments, market prices, network structure, schedule evolution, and equipage of new technologies. AIRLINE-EVOS also models its own heterogeneous population of customer agents that interact with and respond to airline decisions. We describe this model, validation efforts, and a proof-of-concept experiment that demonstrates its capability for assessing policies that balance ATS stakeholder utilities to achieve greater system efficiency, robustness, and safety.

1 INTRODUCTION

The National Aeronautics and Space Administration (NASA) is directly addressing the fundamental research needs of the Next Generation Air Transportation System (NextGen), a substantial and long-term change in the management and operation of the current version of the U.S. ATS, the National Airspace System (NAS). NextGen policies will encompass all airports, airspace, commercial airlines, and other aviation operators under the authority of the Federal Aviation Administration (FAA), the civil aviation authority body of the United States.

In collaboration with NASA, LMI developed a research approach and computational framework to investigate system-wide ATS performance impacts due to ATS stakeholder behaviors—in particular, behaviors of U.S. commercial airlines—under the influences of NextGen and other potential future policies and technology. The ATS is a highly interdependent and complex network of systems and subsystems; operators, regulators, users, and the flying public; policies, procedures, and rules; and facilities and resources. A change in any aspect of this system has cascading effects that ultimately influence the safety, performance, environmental impact, and economics of the ATS as a whole. These stakeholder-level decisions and behaviors have both a tactical and a strategic perspective, and they are influenced by socioeconomics, technology, and policy interactions. A research approach to address this problem must explicitly account for these agent-level behaviors and interactions to provide the detailed insights necessary to guide policy and incentive designs that lead to multidimensional system performance improvements while balancing stakeholder perspectives.

2 RESEARCH APPROACH

In our literature review, we found a number of approaches that have addressed aspects of our research objectives. However, those approaches generally treated individual aspects of our objectives in isolation, rather than comprehensively. A number of analytically based models, such as LMINET (Long et al. 1999), and agent-based ATS-wide simulation models (Sweet et al. 2002; Bilamoria et al. 2000; Volf, Sislak, and Pechoucek 2011), have been developed to assess impacts on the ATS from flight activity, but they do not relate those impacts to airline decisions. Other agent-based modeling approaches address the dynamics of airline decisions, but are limited in scope with respect to number of airlines and/or markets (Kuhn Jr. et al. 2010; Darabi, Mostashari, and Mansouri 2014), do not leverage consequential ATS-wide impacts into airline decisions (Mavris and Garcia 2007), or they account for ATS-wide impacts in their approach but without high-fidelity modeling of aircraft operations and flight interactions in the airspace (Neidringhaus 2004; Gurtner, Valori, and Lillo 2014). Incorporation of such high-fidelity modeling enables greater flexibility for assessing fluid NextGen operational concept scenarios—especially those related to equipage of new technologies—as well as highly interdependent emergent effects such as airspace congestion. To achieve our research objective of supporting NextGen-related decision making, we needed to bridge these individual short-comings with an approach that integrates demand modeling, agent-level modeling of airline behaviors, high-fidelity ATS-wide simulation, and as a new feature, environmental modeling.

We began to address this challenge in Horio et al. (2014), in which we presented a computational framework—called the Air Transportation System Evolutionary Simulation (ATS-EVOS). ATS-EVOS integrated disparate simulation models that focused on specific component features of the ATS, such as travel modal choice and demand generation, airline behavior modeling, and ATS-wide assessments of performance and environmental impacts. Extending that research (Horio et al. 2015), we further developed the computational framework by implementing feedback loops between these modeling components. This feedback enables the dynamics of airline decisions to incorporate measures of ATS-wide impacts, and account for influences related to operational performance, economic outcomes, competitive forces, and environmental impacts by learning from prior decisions and accounting for those experiences in future actions. Figure 1 shows a high-level diagram of the ATS-EVOS framework.

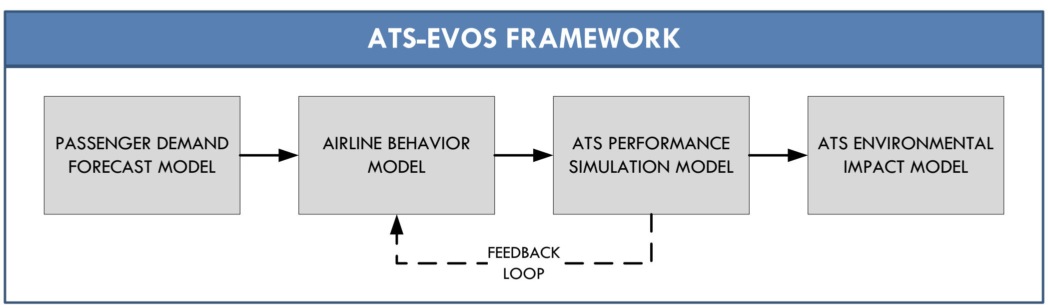

Figure 1: High-level diagram of the ATS-EVOS framework showing how component models are used to comprehensively assess the impacts of airline behaviors in the ATS over time.

This paper presents the improvements and additional progress made in this research, specifically that which is related to the continued development of AIRLINE-EVOS, the airline behavior modeling component of our larger research framework, ATS-EVOS. AIRLINE-EVOS is an agent-based model (ABM) we developed in the Java programming language, designed to enable a better understanding of airline behavior and the functional relationships airlines have to the other ATS stakeholders, both current (airlines, customers, cargo carriers, the FAA) and projected, such as operators of unmanned aircraft systems (UASs). The most significant improvements to AIRLINE-EVOS and our agent-based approach to modeling airline behaviors include:

- New learning functionality that allows AIRLINE-EVOS to run the entire computational framework of models over multiple iterations, allowing the state of the system and agent behaviors to be influenced by prior states and evolve over time.
- Airline agents have new strategic options available, including modifying departure time, and modifying their service network structure. These real-world strategies in tandem additionally allow airlines to augment their frequency of service to particular markets, and provide viable responses to certain new policy implementations, such as congestion pricing at airports.
- Customer agents have an estimated parameter for willingness-to-pay (WTP), allowing for utility maximizing ticket purchase decisions to be more representative of real world behavior.
- AIRLINE-EVOS—specifically airline pricing parameters and customer WTPs—has been calibrated to NASA demand forecasts, making the instantiation of customer agents in the model more representative for evaluations of potential policy impacts in projected future year scenarios.
- AIRLINE-EVOS results after calibration were validated against empirical airfare data collected by the Bureau of Transportation Statistics (BTS), with promising results.

3 AIRLINE-EVOS MODEL DESCRIPTION

The following subsections provide a high-level description of the different agents, and of the behavioral mechanics of AIRLINE-EVOS. A more complete model description following the Overview, Design concepts, and Details (ODD) protocol (Grimm et. al. 2010), may be referred to in the appendix of Horio et al. (2015).

AIRLINE-EVOS models two types of entities within the system: airlines and customers. The model purpose is focused on the exploration of the dynamic interactions between customer ticket choices, airline decisions, and the performance of the ATS as a whole. In its current development, AIRLINE-EVOS has modeled 26 different airlines—which represent three different business model categories for legacy or network carriers, low-cost carriers, and regional carriers—and their interactions with approximately 950,000 individual customers. Airline agents in a current-day schedule scenario manage 22,000 daily flights and offer airfares and itinerary offerings for over 54,000 potential O-D pairs. The model is able to exceed these parameters, however we have not fully explored its computational limits and cannot therefore specify parameter limits.

3.1 Airline Agents

Airlines have the explicit objective of maximizing short-term profit. We identified two key behaviors that airlines exhibit towards this objective, in response to policy changes and new technologies: airfare price manipulation, and schedule modifications. Airline agents make 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; the agents gradually learn the ideal airfare for a particular market, and allow the airfare to drift when significant changes to the market affect the pattern of customer behavior. These dynamic pricing changes are based on customer advance purchase time characteristics and the available seats for a given flight, and give rise to several emergent behaviors in AIRLINE-EVOS.

- Market-based dynamics of customers, with respect to purchasing behaviors over time, and travel trends by traveler type and origin-destination O-D pair (e.g., leisure travel trending towards lower volume in certain markets).
- Competitive airline behaviors, with respect to how airlines change airfares over time, in the subsequent impact on profit, and the resulting evolution of market share.

- Airline operations, with respect to load factor trends and equipment gauging by market, in addition to decisions to equip new technology.
- Other emergent effects at the ATS-level are expected to result after running AIRLINE-EVOS schedule output in the ATS-wide simulation component of our larger research framework. These operations-related performance metrics include measures of airspace congestion, delay, violation of safe separation, and others.

Airlines also adjust the flight schedule by changing their equipment assignments, modifying their network structure, or changing departure times. Scheduling changes by the airline agents are strategic decisions that attempt to increase revenues by improving captured market share, and by better matching capacity with demand.

These airline behaviors also take into account feedback they receive from operational models that simulate system-wide impacts from airline decisions. The feedback loop effectively allows airlines to observe the eventual success or failure of their attempts to capture greater profit through the individual decisions they make, and enables their actions to evolve over time, based on what they observe. For example, one particular airline decision may be revealed (through simulation) to be suboptimal due to unexpected congestion delays, whose costs negate any benefit expected prior to simulation. Emergent system impacts that result from airlines simultaneously making independent—but highly interconnected—decisions, may also generate unexpected effects that might influence future decisions. This is one of the modeling areas in which the most progress was made since Horio et al. (2014).

3.2 Customer Agents

In every iteration of the model, customers explicitly choose the ticket that will maximize their utility, taking into account their individual preferences, including market-specific WTPs. The traveling public is the only source of revenue in the system. This is of significant importance because their ticket purchase decisions serve as triggers for airline behaviors. Customer agents are modeled as heterogeneous agent populations, with a number of differentiating attributes, including the O-D pair that defines the agent's desired route. In AIRLINE-EVOS, their behavior is strictly concerned with making a decision about whether or not to purchase an airline ticket, and if so, which ticket to purchase. They select which available ticket they will purchase, based on a cost- and inconvenience-minimizing utility function, with some degree of randomness. We assume that customer agents might be strongly influenced by other factors, such as loyalty programs, and therefore with some probability, do not strictly maximize utility decisions. After selecting airline tickets, the role of the customer agents in the model has been satisfied.

3.3 Model Temporal Scale

The spatial scale of the model is limited to the measure of great circle distances between origin and destination points in nautical miles, which is used for calculating fuel burn, travel duration, and operating costs.

Temporal scales in the model are dimensionless; explicit representation of a fixed time unit for each simulation cycle is not specified. Each iteration of AIRLINE-EVOS is reflective of an adaptive process in which airlines will assess their performance and make appropriate strategic changes to their airfare pricing and schedules in response to changes in the competitive environment. Competing airlines also make adjustments, to the effect that when they are responding to customer demand patterns, they are also responding to the competitive actions of other airlines. In an iterative way, the market is reengaged by all airlines after any changes, and the airlines determine success or failure of those changes with respect to profitability. The outcome of this assessment influences subsequent strategy decisions by the airline. Each learning iteration loop is considered the next available time for the airline to publish and implement a flight schedule change; it is not a specified increment of time. It is also assumed that all learning iterations and the resulting model outcomes for experiment scenarios are within context of the same seasonal period of the starting schedule.

3.4 Process Overview and Scheduling

The following conceptual diagram in Figure 2, shows the mechanics, and related inputs and outputs for the AIRLINE-EVOS model.

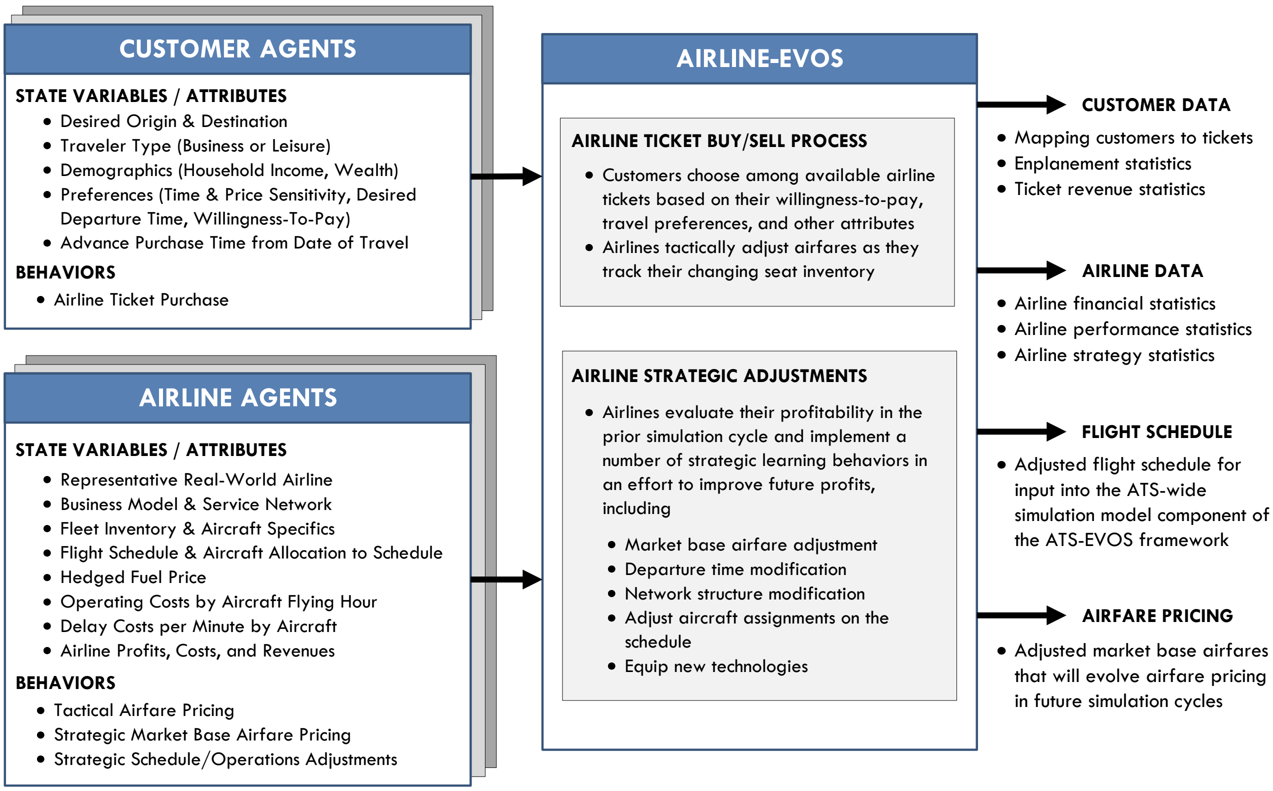

Figure 2: Conceptual model diagram for AIRLINE-EVOS.

Following the outline in Figure 2, a customer agent population is generated for each O-D pair based on market elasticity assumptions, calibrated (Kumar et al. 2015) to historical ticket sales data (Department of Transportation 2015), and to experiment scenario results from the Transportation Systems Analysis Model (TSAM) (Baik and Trani 2005), a national demand forecasting model. This results in a calibrated WTP for each customer. Customers are also assigned other attributes based on demographics from TSAM, and other assumptions such as how far in advance the purchase is from the flight departure date.

An airline agent population is also instantiated, based on predetermined representative airline business models. Each airline agent loads an initial flight schedule, aircraft equipment-related data, and fleet allocation assignments. Calibration of AIRLINE-EVOS to TSAM also updates an airline's base market price.

AIRLINE-EVOS at a high-level, performs two functions in which (1) airlines sell tickets to customers, and (2) airlines conduct strategic behaviors to improve profits.

- **Airline Ticket Buy/Sell Process**: Airlines generate feasible customer itineraries from their initial airline schedules, and price those itineraries according to their own airline-specific costs, and the base airfares for a given market. These priced itineraries are offered to individual customers.

Customer agents are ordered by their advance purchase time, with earlier purchase dates getting first choice of tickets. Customer agents make a utility maximizing choice—with some accommodation for irrationality—to determine which airline ticket to purchase, considering only those tickets under their own WTP. Tickets are sold one at a time, and as tickets are sold, airlines track their seat inventory and profits. Airfares are tactically adjusted with each ticket transaction, until all customers are served.

- **Airline Strategic Actions**: After all customer agents have made their selection, airline agents assess their current operational and financial state. Given strategic behavior rule-sets that consider individual markets, airlines initiate or respond to competitive actions or adopt a different economic operating point, represented by a base starting airfare in that market. Equipage decisions for new technology are also made at this point in the AIRLINE-EVOS process. These strategies are applied in a sequential order—based on SME input—and include (1) market base airfare adjustment, (2) network modification, (3) departure time modification, and (4) equipment/gauge swapping. For each type of airline strategy implemented, a feedback loop is executed in which the same customer agent population is reengaged to consider the resulting new set of itinerary and airfare options. This is representative of an airline trying to improve its performance by testing new schedule and airfare strategies at the next available incremental schedule change opportunity. It is a "learning" process, repeated until some convergence or modeling threshold is reached. These strategies are all part of a single iteration of AIRLINE-EVOS. Using the ATS-EVOS framework, AIRLINE-EVOS can then take in as inputs ATS-wide impacts as a results of AIRLINE-EVOS outputs in the previous cycle. The end results is multiple iterations of AIRLINE-EVOS within each iteration of ATS-EVOS. These nested iterations of AIRLINE-EVOS helps facilitate a convergence of strategic decision outcomes and further mitigate any undesirable effects due to the specific ordering of strategies.

AIRLINE-EVOS results in four generalized outputs. The primary output is an adjusted flight schedule; all the model interactions between customers, airlines, and the system result in a modified flight schedule that may then be input into an ATS-wide simulation as part of the larger ATS-EVOS framework. The resulting flight schedule aggregates the adjusted flight schedules from all the modeled airline agents. The model also aggregates customer and airline data, tracking enplanements, revenues, costs, and profits. The resulting airline decisions for schedule and airfare pricing, along with the associated ATS-wide simulated impacts, are used as inputs into the next iteration of AIRLINE-EVOS, enabling an evolution of the system that accounts for learning from prior states.

3.5 Customer Ticket Choice Utility Function

AIRLINE-EVOS is a detailed model with a number of submodels that are used for the full range of agent behaviors from ticket choice to airline strategic decisions and assessment of financial performance. One submodel that is worth discussing here, is the ticket purchasing utility function used by the customer agents, primarily because it is a significant driver for the results of our validation and proof-of-concept experiment described later in this paper.

The process revolves around customers choosing the ticket that will minimize their disutility (i.e. maximize their utility), although they only consider a subset of tickets they have identified with airfare amounts below their WTP. Customers seek to balance airfare amount while also choosing an itinerary that offers the least inconvenience in terms of its total travel time and desired departure time. Given calculated utilities for all valid ticket options, customers determine if they will rationally choose the ticket that would best satisfy their utility and preferences, or whether they will instead irrationally choose an alternative with less than optimal utility. Irrational alternative choice probabilities are proportional to the distance (i.e., difference in utility) between the most preferable ticket and the worst alternative.

For customers making rational ticket choices, each itinerary i is evaluated by a customer agent g, making a determination of the utility value U_{gi}, as given in (1). This formulation was adapted from research by Mavris and Garcia (2007), with our addition of a scaling term to allow for weighted customer travel preferences and also for the normalization among ticket options in the equation terms.

$$U_{gi} = -\varphi_{gi}\left[C_{g,o-d}\left(\frac{Fare_i}{Fare_{MIN_i}}\right) + S\left(T_g\left(\frac{FlyHr_i}{FlyHr_{MIN_i}}\right)\right)\right]. \tag{1}$$

This utility function is a function of the ticket's associated airfare, the duration of travel associated with itinerary i, and the difference in the itinerary departure time from the desired departure time. The terms inside the brackets of the utility function account for the airfare and the duration of travel. The first set of these terms in the brackets, normalizes the fare $Fare_i$ being evaluated against the lowest fare available for travel between the desired O-D pair, $Fare_{MINi}$. The airfare sensitivity coefficient $C_{g,o-d}$ is specific to the traveler type of customer g and to the O-D market for itinerary i, and is a measure of the importance customer g places on fare in determining the best overall ticket.

The second set of terms inside the brackets, concerns the total duration of travel for itinerary i. Total travel duration $FlyHr_i$ is normalized against the shortest travel duration of all itineraries of the tickets being considered, $FlyHr_{MINi}$. The value of time coefficient T_g is specific to customer g and is estimated from the household income attribute assigned to the customer at instantiation. It reflects the importance customer g places upon total travel time in determining the best overall ticket. The S term is a scaling factor that equalizes the difference between the different units for price and duration used within the utility brackets. The outside term φ_{gi} acts as a penalty modifier in the calculation of itinerary disutility that accounts for any difference in desired arrival time and actual arrival time, used in the same way as Mavris and Garcia (2007). In general, the more the arrival time of i deviates from the desired arrival time of customer g, the harsher the penalty.

4 MODEL VALIDATION

A primary component of our validation efforts was quantitative, and the most accessible and relevant data we have for comparison in this respect, relates to airfares. The airfare database we have chosen to use is the BTS Airline Origin and Destination Survey (DB1B) (Department of Transportation 2015), which is a 10-percent quarterly sample of all tickets sold by U.S. commercial airlines. The distribution of DB1B trends also provide a validation pattern for the entire ticket purchasing process, including how airlines price tickets and how customers, in turn, decide which tickets to buy, if at all. The validation of this overall process effectively validates AIRLINE-EVOS in its entirety, with respect to sufficiently capturing the complexity of the internal system mechanics to ultimately result in overall system outcomes that are representative of the real world. Figure 3 shows simulated results from AIRLINE-EVOS and its comparison to relevant DB1B results.

We do not expect the comparison of results in Figure 3 to exactly match, because we are simulating only a single representative day and are modeling simplified airline behaviors that may not perfectly capture factors inherent in the DB1B data, including demand seasonality, detailed accounting for fare classes and airline yield management practices, and day-of-week pricing strategies. We do seek to use this validation pattern to help determine if our abstraction of airfare pricing logic, which considers only when only considering advance purchase day, load factor at time of purchase, market-specific base airfare as a result of customer WTP, and options for nonstop or a single connection itinerary, sufficiently captures the complexity of real-world pricing. Given these considerations, comparison of mean airfares in Figure 3 suggests that AIRLINE-EVOS airfare pricing logic is fairly representative of the real world. Customer decisions about tickets to purchase are dependent on their price, thus validation of consumer choice—as observed by the comparison of distribution shape between AIRLINE-EVOS and DB1B—simultaneously offers concurrent validation for the ticket purchasing utility decisions of our customer agents.

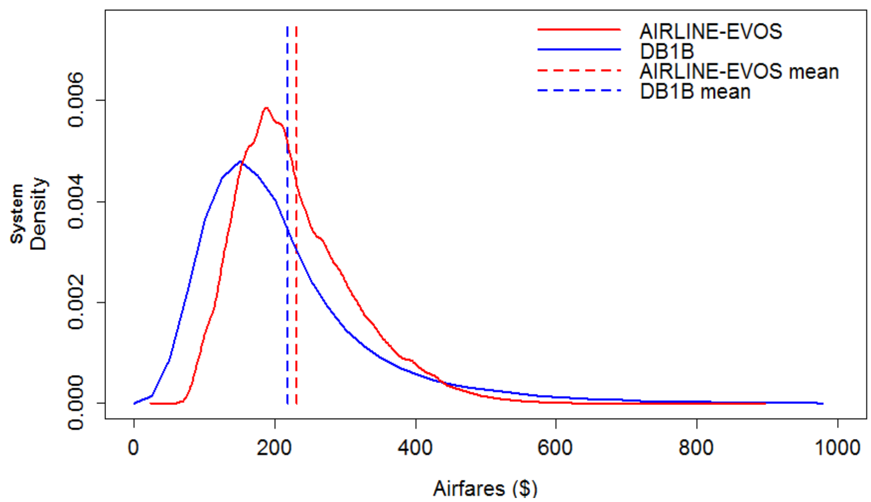

Figure 3: Simulated airline airfares in comparison to empirical airfare data.

5 PROOF-OF-CONCEPT EXPERIMENT

We present here a proof-of-concept experiment and our results. This effort was to provide us confidence that our research approach and modeling framework, when used in rigorous analytical studies, would yield decision-quality results. These results however are not designed for, nor intended to directly inform policy decisions. Future analyses would follow a similar experimentation process but would require more complex and detailed experiment parameters.

The experiment we describe here is for assessing system impacts due to airline response to a notional carbon tax. There is an increasing recognition and acceptance of the fact that the environmental cost of growth and progress should be minimized where possible. One of the ways in which Governments across the world sharply reduce carbon emissions (i.e. environmental cost) is by imposing carbon tax. A carbon tax is an "upstream" tax on the carbon content of fossil fuels (coal, oil, and natural gas) and biofuels. We use an potential implementation of a carbon tax as the experiment scenario in which airlines must adapt their responses to the increase in operating costs.

This particular experiment was chosen for its potential demonstration for representative modeling of airline behaviors, in particular, adaptive response to conditions of higher operating costs. It is our opinion that demonstrating representative mechanics for associated behavioral responses and showing that emergent outcomes are plausible and reasonable, provides strong evidence that this model may be used for research into many topics of interest for NASA and the larger aviation research community.

On the basis of our previous research in Horio et al. (2014) and input from industry subject-matter experts (SMEs) on the research team, the most likely real-world implementation of a carbon tax would be as a fuel surcharge. The objective of the tax would be to create an incentive for airlines to decrease their contribution to environmental pollution, or otherwise encourage airlines to burn less fuel. Technically, this would mean a fee levied on the amount that airlines consume. An equivalent yet much simpler implementation is to levy the additional cost on the fuel that airlines purchase. For our purposes, this is equivalent to an increase in fuel price. Airlines will pay their own normal fuel price, plus the additional amount levied by the civil aviation authority as a carbon tax. The policy also seeks to help prevent airline overscheduling behaviors by encouraging upgauging—the reallocation of equipment to use larger aircraft in profitable markets and better absorb operating costs—assuming that a larger aircraft requires less fuel per passenger mile.

The experiment represents a simulation of the 2012 airline market and schedule in which the U.S. airline industry was operating. We model fuel price increases in AIRLINE-EVOS through the use of a price multiplier. Our experiments assume a notional carbon tax multiplier of 2.0, which, in effect, will maintain differentiated airline fuel prices due to hedging—the strategy by which airlines buy future fuel at current prices in the hopes of achieving a price advantage—but globally doubles all fuel unit costs for airlines in the model. The high tax amount is partly to help exaggerate outcomes and identify the extent to which changes in AIRLINE-EVOS are observed. This experiment will illustrate how airline agents sense revenue implications and adapt their pricing and scheduling responses to maintain profitability.

6 RESULTS

The carbon tax policy is designed to limit the environmental cost of growth in air travel demand by encouraging airlines to make "greener" business decisions. The ATS-EVOS modeling framework is designed to provide quantitative aid to decision makers/policymakers in understanding the effect of such a policy on the operators (airlines) and the users (customers). Parameter values should exist such that a carbon tax policy will force airlines to change their operations by modifying schedule or equipment assignments. The higher the carbon tax rate, the more airlines will change their operations.

Using our calibrated airline parameters for airfare pricing and customer WTP values, we simulated three iterations of AIRLINE-EVOS as part of the ATS-EVOS computational framework that integrated additional NASA simulation models. We used a baseline schedule that represented the year 2012. The computational cost for a full iteration of ATS-EVOS was such that we could only afford three iterations for this proof-of-concept. We ran the three iterations for a baseline case and for the carbon tax scenario; additional iterations were desired however our selected ATS-wide simulation component in the ATS-EVOS loop had an extremely high computational cost and required 12+ hours per run for just that part of the simulation. We then explored the results for impacts of the policy on airfares. Figure 4 shows the resulting mean airfares for the system, along with how airfares changed with the implementation of the five strategies available to the airlines for adapting to system and operational changes. Figure 5 shows corresponding system enplanements—how many customers bought tickets and boarded an aircraft—for the baseline and carbon tax scenario.

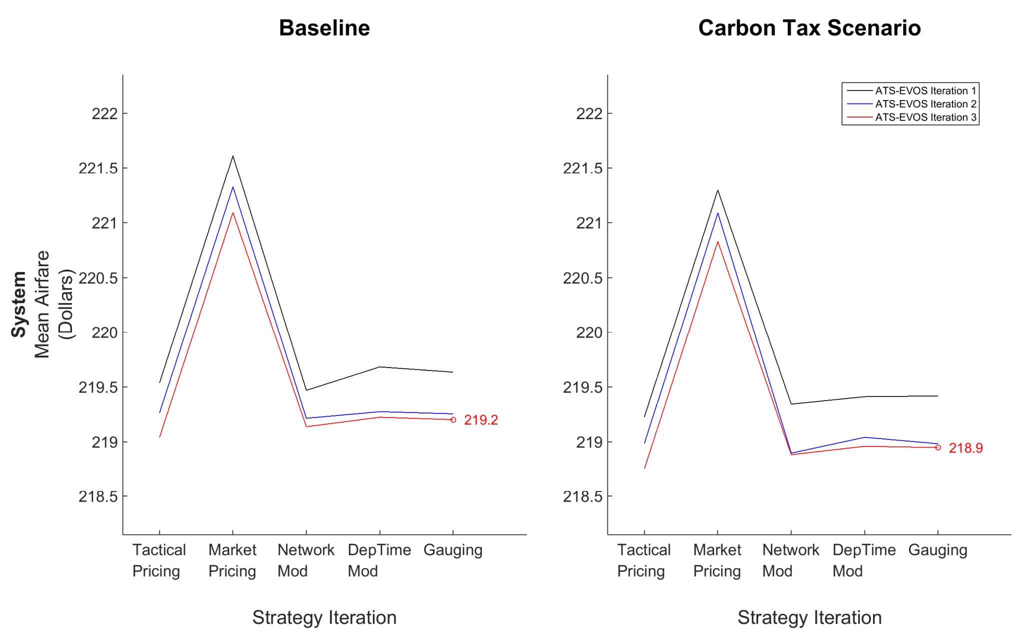

Figure 4: System Mean Airfares for 2012 Baseline and Carbon Tax Scenario.

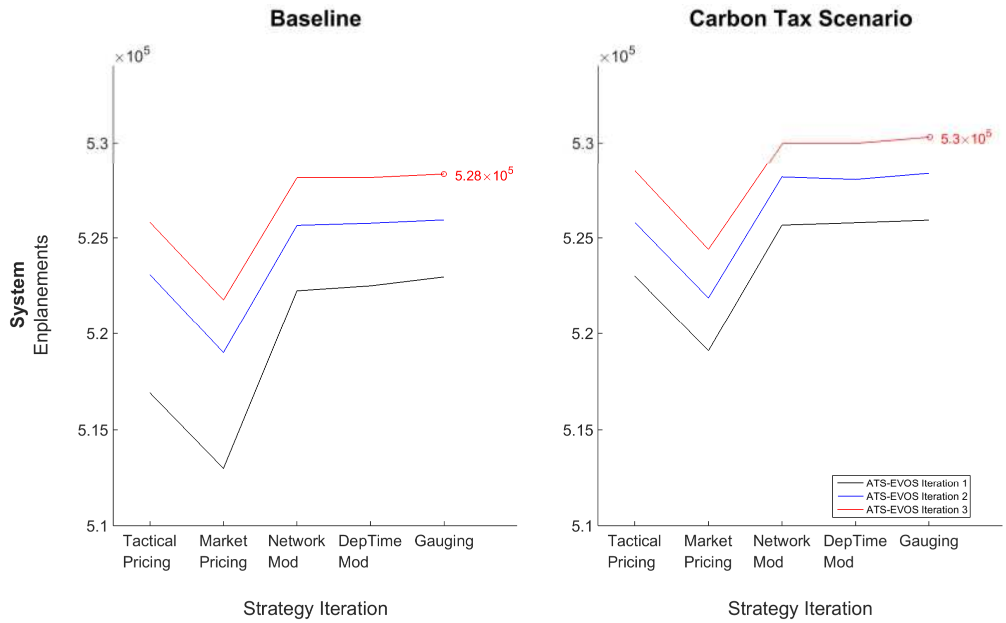

Figure 5: shows enplanements for the system for both the baseline and experiment case. Results indicate that the carbon tax policy reduced the variation in results between iterations, and increased enplanements for the experiment case. Figure 2 shows that the system-wide airfare in the carbon tax case trended lower than in the baseline case, resulting in an increase in enplanements.

Each of the plots includes one line for each of the different ATS-EVOS iterations to provide insight into the evolution of the model. The x-axis categories represent our five airline behaviors for tactical and strategic response. These strategies were executed in sequential order, from left to right, thus reading the plot in that direction indicates model progress within an ATS-EVOS iteration, and reading the plot with comparison between the different ATS-EVOS iteration plot lines indicates model evolution over the ATS-EVOS iterations. One way of interpreting these plots is to observe the last data point for each iteration, because it represents the result of all strategies and learning for that iteration. Observing the last data point for the latest ATS-EVOS iteration—highlighted by a red data point and value label—is one way of interpreting the most emergent state for the simulated results.

In this experiment, we observed that a carbon tax makes flights more expensive, and airline agents respond strategically to mitigate those penalties, adjusting their operations with respect to their schedule and base airfares at individual markets. As these mitigating strategies accumulate, the resulting emergent state is a reduction in airfares, which leads to an increase in enplanements. The overall reduction in airfares also makes connecting flights more attractive—as they are typically cheaper—so consumer choice shifts away from nonstop flights. Consequential impacts include an upward trend in system delays, potentially as a result of schedule adjustments, although the overall magnitude of the delay is negligible.

As expected, the operational costs for the airlines increased significantly throughout the system. Profits correspondingly decreased significantly. Contributing to the loss in profits was an unexpected decrease in system-wide average airfares under the carbon tax policy, while at the same time costs were increasing. Further evidence of this needs to be corroborated in future research.

While some of the results we've described above are logically consistent, such as lower airfares corresponding with higher enplanements, the overall results are counterintuitive. A carbon tax would increase airline costs, and when operating costs increase, real-world airlines typically (1) raise airfares and pass on costs to customers, causing a decrease in enplanements, and (2) cancel non-profitable flights, or (3) a combination of both. This outcome has highlighted the need for some enhancements to our current model, to allow airline agents to act in more realistic ways in response to operational cost increases. Airlines in our model have variable schedule costs, but with a fixed minimum value—as a result of not being able currently to cancel flights or remove aircraft from service—and as a result, can only respond by working to aggressively try and sell more seats. Given the static nature of their WTP values, airlines are forced to drop airfares to capture as much revenue as possible to offset the operational cost increases imposed by the carbon tax. Future research may address these issues by providing airline agents with the strategic option to remove aircraft from service. These results, though counterintuitive, are explainable, reasonable, and provide a clear direction for future enhancements to our model. They also highlight the challenge in developing a proof of concept model of the commercial aviation system, which is complex enough to limit the usefulness of a simplified model.

7 CONCLUSIONS

To substantiate our research we conducted a set of proof-of-concept experiments for demonstrating the capability of our model and our research approach for assessing policies that balance ATS stakeholder utilities to achieve greater efficiency, robustness, and safety in the ATS.

Overall, the results of our final experiments showed only relatively small changes with respect to the baseline. These experiments were successful in that we demonstrated the ATS-EVOS approach and our AIRLINE-EVOS model of airline behavior as viable tools that can produce meaningful metrics for analyzing airline decisions and the consequences of those decisions. Our observed results, while modest, exhibit directional trends that seem to be reasonable, explainable, and representative of the real-world system. We also validated the results with our team of industry experts. The results from our experiments are not designed or intended to directly inform policy decisions. Future analyses would use the experimentation process outlined in this report but would require more complex and detailed experiment parameters, and some modeling enhancements.

8 FUTURE RESEARCH

Work is underway to address some of the shortcomings discussed in our results. We are exploring (1) the inclusion of strategies for removing a flight from service, (2) augmenting market-based WTPs to more accurately represent a carbon tax scenario, and (3) exploring preferential weights for ticket choice utility.

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 thank Jerry Smith and Nelson Guerreiro from NASA, who directly supported this NRA. We also thank William Cotton of Cotton Aviation Enterprises and Vince Costanzo for their counsel and advice, and Jeffrey Goldsmith, James Hebden, Shahab Hasan, Virginia Stouffer, and Terry Thompson of LMI, for their contributions and support of this research. Thank you also to the anonymous reviewers for their help.

REFERENCES

Baik, H., and A. A. Trani. 2005. "A Transportation Systems Analysis Model (TSAM) to Study the Impact of the Small Aircraft Transportation System (SATS)." In *Proceedings of the 23rd International Conference of the System Dynamics Society*.

Bilamoria, K., B. Sridhar, G. Chatterji, K. Sheth, S. Grabbe. 2000. "FACET: Future ATM Concepts Evaluation Tool." In *Proceedings of the 3rd USA/Europe Air Traffic Management R&D Seminar*.

Department of Transportation, Bureau of Transportation Statistics. 2015. “DB1B Database.” Accessed April 8, 2015. http://www.transtats.bts.gov/Tables.asp?DB_ID=125.

Darabi, H. R., A. Mostashari, and M. Mansouri. 2014. “Modelling Competition and Collaboration in the Airline Industry using Agent-Based Simulation.” *International Journal of Industrial and Systems Engineering* 16(1): 30–50.

Grimm, V., U. Berger, D. L. DeAngelis, J. G. Polhill, J. Giske, and S. F. Railsback. 2010. “The ODD Protocol: A Review and First Update.” *Ecological Modelling* 221(23):2760–2768.

Horio, B. M., A. H. DeCicco, V. L. Stouffer, S. Hasan, R. L. Rosenbaum, and J. C. Smith. 2014. “Agent Based Modeling of Air Carrier Behavior for Evaluation of Technology Equipage and Adoption.” In *Proceedings of the 14TH AIAA Aviation Technology, Integration, and Operations Conference*. Reston, Virginia: American Institute of Aeronautics and Astronautics.

Horio, B. M., V. Kumar, A. H. DeCicco, S. Hasan, V. Stouffer, J. C. Smith, and N. M. Guerreiro. 2015. “An Integrated Framework for Modeling Air Carrier Behavior, Policy, and Impacts in the U.S. Air Transportation System.” In *Proceedings of the 15TH AIAA Aviation Technology, Integration, and Operations Conference*. Reston, Virginia: American Institute of Aeronautics and Astronautics.

Kuhn Jr., J. R., J. F. Courtney, B. Morris, and E. R. Tatara. 2010. “Agent-based Analysis And Simulation Of The Consumer Airline Market Share for Frontier Airlines.” *Knowledge-Based Systems* 23:875–882.

Kumar, V., Horio, B. M., A. H. DeCicco, S. Hasan, V. L. Stouffer, J. C. Smith, and N. Guerreiro. 2015. “Understanding Air Transportation Market Dynamics using a Search Algorithm for Calibrating Travel Demand and Price.” *In Proceedings of the 15TH AIAA Aviation Technology, Integration, and Operations Conference*. Reston, Virginia: American Institute of Aeronautics and Astronautics.

Long, D., D. A. Lee, J. P. Johnson, E. M. Gaier, and P. F. Kostiuk. 1999. “Modeling Air Traffic Management Technologies With A Queueing Network of the National Airspace System.” NASA Contractor Rep. 208988, Washington, DC.

Mavris, D. N., and E. Garcia. 2007. “Decision Making Support for NASA’s Aeronautics Research Mission Directorate.” NAS1-02117, Contract 6075.

Niedringhaus, W. P. 2004. “The JET-WISE Model of National Airspace System Evolution.” *SIMULATION* 80(1): 45-58.

Sweet, D., V. Manikonda, J. Aronson, K. Roth, and M. Blake. 2002. “Fast-Time Simulation System for Analysis of Advanced Air Transportation Concepts.” *In Proceedings of the AIAA Modeling and Simulation Technologies Conference and Exhibit*. Reston, Virginia: American Institute of Aeronautics and Astronautics.

Volf. P., D. Šišlák, and M. Pěchouček. 2011. “Large-scale High Fidelity Agent-Based Simulation in Air Traffic Domain.” *Cybernetics and Systems: An International Journal* 42(7): 502-525.

AUTHOR BIOGRAPHIES

BRANT M. HORIO is an operations researcher with LMI. He holds an M.S. degree in Operations Research from George Mason University. His email address is bhorio@lmi.org.

VIVEK KUMAR is a research engineer with LMI. He holds a Ph.D. and M.S. degrees in Systems Engineering and Operations Research from George Mason University. His email address is vkumar@lmi.org.

ANTHONY H. DECICCO is an Operations Research Analyst with the RAND Corporation. He holds an M.S. degree in Operations Research from George Mason University. His email is adecicco@rand.org.