Right of Way in the Sky: Two Problems in Aircraft Self-Separation and the Auction-Based Solution

Ido Erev, Technion – Israel Institute of Technology, Haifa, Israel, **Greg Barron,** Harvard Business School, Boston, Massachusetts, and **Roger Remington,** NASA-Ames Research Center, Moffett Field, California

There has been a growing movement to give commercial airliners more freedom in choosing their routes and responsibility for detecting and avoiding conflicts. These "free flight" concepts must contain new rules for assigning right of way in potential conflict situations. To evaluate the effect of prospective rules, the current paper derives the expected response of agents who exhibit different levels of sophistication. Traditional game theoretic analysis is used to derive the behavior of rational agents. Computer simulations are used to predict the behavior of boundedly rational reinforcement learners. The results reveal that several seemingly reasonable, straightforward right-of-way rules might lead to undesirable outcomes. These problematic results are robust to the assumed level of rationality. It is shown that these problems can be alleviated by using auctions to resolve competition for right of way. Actual or potential applications of this research include the usage of second price auctions to address right-of-way and similar conflicts.

INTRODUCTION

Recent developments in aviation technology significantly improve the information available to pilots and the possibility of communication between aircraft (e.g., Livack, McDaniel, Battiste, & Johnson, 1999). In the near future, equipment will be available commercially that will allow aircraft to connect to air traffic control central computers and receive an updated picture of their location and the air traffic around them. The expected availability of this technology has led the U.S. Federal Aviation Administration (FAA) to consider changing current flight regulations to take advantage of the potential increases in efficiency afforded by this new technology (see National Civil Aviation Review Commission, 1997; Radio Technical Commission for Aeronautics [RTCA], 1994; Wickens, Mavor, & McGee, 1997). Among the options being investigated are changes to the rules that determine who is responsible for maintaining adequate separation between aircraft and who should maneuver to avoid potential conflicts (RTCA, 1994).

In current air traffic operations, ground-based air traffic controllers have the responsibility for maintaining separation as well as the authority to vector aircraft to avoid separation violations. Pilots fly in accordance with air traffic control instructions regarding airspeed, altitude, and heading, unless such instructions pose a direct threat to the aircraft. In effect, air traffic control acts as a noncompetitive entity within the competitive air carrier environment, in which the principal concern is the safe operation of the national airspace. The changes under consideration would replace the centralized oversight provided by air traffic control with a decentralized decision-making system. Individual aircraft could freely choose fuel- or time-optimized routes, instead of being routed along fixed airways, as is the typical case today. This general concept has been termed *free flight* and is the subject of intensive investigation by the FAA, the National Aeronautics and Space Administration (NASA), and various industry and academic research efforts (e.g., the Advanced Air Transportation Technologies program).

Address correspondence to Ido Erev, Faculty of Industrial Engineering and Management, Technion, Haifa 32000, Israel; erev@tx.technion.ac.il. **HUMAN FACTORS,** Vol. 46, No. 2, Summer 2004, pp. xxx–xxx.

One change envisioned under free flight is the transfer of responsibility for separating aircraft from ground-based air traffic control to individual aircraft, a concept referred to as *self-separation.* Under self-separation, aircrews would be alerted to potential conflicts and be responsible for choosing maneuvers that would resolve them. Two implementation approaches have been investigated: (a) communication between conflict aircraft to negotiate a solution maneuver (e.g., Livack et al., 1999; Lozito et al., 2000) and (b) "right-of-way" rules that would determine which aircraft should maneuver. Proposed right-of-way rules rely on existing regulations in effect today under visual flight rule conditions. Visual flight rules are in effect for flight operations outside of terminal control areas conducted under visual conditions that afford good visibility. These rules state that when two aircraft are in potential conflict (i.e., continuing on the current trajectories would lead to a separation violation), the right of way goes to (a) the aircraft that comes from the right or, in cases of overtaking, (b) the aircraft in front. Whereas these right-of-way rules have been found to be very useful for many years, their effectiveness when applied to the entire airspace has yet to be demonstrated. Likewise, negotiating an avoidance maneuver sounds reasonable but is largely untried.

The development of new, efficient rules requires good understanding of the technological and economic constraints and of the expected behavior of future users. It is of particular importance to try to understand the possible long-term adaptation of users to the new rules. Failure to do so can lead to unintended consequences that may be inefficient or potentially dangerous. However, existing methods make it difficult to assess long-term adaptation. Performance under specific self-separation rules has been studied in large-scale human-in-the-loop simulations (e.g., Endsley, Mogford, Allendoerfer, Snyder, & Stein, 1997; Fleming, Lane, & Corker, 1999; Galster, Duley, Masalonis, & Parasuraman, 1998; Hilburn, Jorna, Byrne, & Parasuraman, 1997; Hoffman, Zeghal, Cloerec, Grimaud, & Nicolaon, 2000; Lozito et al., 2000). These studies have looked at performance measures such as error, conflict detection times, and overall reports of workload. Such empirical studies are expensive and, because of their short duration, cannot provide insight into adaptive strategies that emerge with long-term exposure. Indeed, though the effects of competition on airline scheduling and information dissemination practices have been recognized for some time, scant attention has been directed at the possible effects of competition for flight routes that may result from self-separation.

The main goal of the current research is to highlight the behavioral implications of some options considered by the designers of the new rules. Our approach was to use a method that would allow examination of a large space of options without the need for detailed scenario specification. Indeed, human factors practitioners have long attempted to establish a stronger voice early in concept development. However, it has been difficult to provide a rigorous analysis at that early stage because many human factors techniques require extensive specification of interface, tasks, and other aspects of the context that are available only late in design. For example, large-scale simulation must commit to detailed equipment, procedures, and scenarios. Without good simulation of the relevant environment, is it hard to derive accurate prediction of the likely behavior.

The current attempt to address this difficulty rests on two observations. First, traditional game theoretic analysis can be used to derive the expected behavior of rational agents. Second, when human agents can learn from experience, the direction and the magnitude of the deviations from rational choice can be captured with models of reinforcement learning (see Erev & Roth, 1998). These observations suggest that derivation of the behavior of rational agents and of reinforcement learners implies a prediction of the range of expected behavior. When the range is narrow, the two models converge on common behavioral expectations. In such cases, the two-model analysis can be used to highlight the behavioral implications of different rules.

As a corollary to this goal, we also explore game theoretic analyses as a means of developing more efficient rules. To maximize efficiency, the new rule should be sensitive to the aircraft's ability to change course and to the relevant costs and incentives. For example, if one of the aircraft is about to miss an important connecting flight, it seems reasonable that it should be preferred. Whereas the present research focuses on right of way in the sky, we hope that the analysis at hand will also shed light on the value and limitations of the two levels of rationality analysis we use to address other regulations.

METHOD

An attempt to derive the implications of game theoretic analysis for the current tasks reveals that it is not clear how the natural problems of interest can be abstracted as decision problems. Game theoretic research focuses on the factors that affect decisions in simple problems, but it provides only limited help in simplifying complex problems. Given that there is no good solution to this problem, the current research does not try to study "natural" problems. Rather, we chose to focus on simple scenarios that demonstrate interesting dilemmas. In particular, we start by focusing on simple situations in which reasonable right-of-way rules create inefficiencies. We continue by proposing solutions to these dilemmas. By abstracting the problem, our analysis highlights general characteristics of encounters that may apply across a wide range of traffic conditions. We hope that understanding these characteristics can be useful in future analysis of more complex tasks. For example, in the current analyses we ignored the role of specific systems, such as the Traffic Collision Avoidance System (TCAS). Nevertheless, we believe that our results can be used to suggest variants of TCAS that facilitate efficiency. We return to this point later.

As explained earlier, the current analysis tries to predict choice behavior by deriving the behavior of rational agents and the expected deviations from this prediction. The behavior of rational agents is derived using the assumption (typically made in economic analyses) of selfish agents that maximize expected return, assuming that all other agents do the same. The expected deviations from rationality are derived using computer simulations in which adaptive agents behave according to the quantification of the law of effect (Thorndike, 1898) proposed in Erev and Barron's (2003) model of reinforcement learning. This model is preferred over alternative models of bounded rationality (e.g., Camerer & Ho, 1999; Chen & Tang, 1998; Cheung & Friedman, 1998; Cooper, Garvin, & Kagel, 1997; Daniel, Seale & Rapoport, 1998; Roth & Erev, 1995; Sarin & Vahid, 2001; Stahl, 1999) because it assumes very low level of rationality and provides good approximation of human behavior in a wide set of experimental tasks with a single set of parameters. (Best fit is obtained when the experimental participants have no prior information about the payoff rule. When participants receive prior information, behavior tends to move toward the rational predictions.) The specific assumptions made by this model are presented in the Appendix.

TWO SHORTCOMINGS OF THE CURRENT REGULATION

As noted, current regulations grant right of way to traffic on the right. Whereas this regulation has been rather effective in addressing rare conflicts that are not addressed by ground control, such a position-based scheme is associated with two inefficiencies. To highlight these inefficiencies, we define *100% efficiency two-aircraft conflict resolution* with two properties: (a) The right of way is given to the aircraft that values it more, and (b) the conflict does not incur additional costs on the system. Specifically, if the cost from losing the right of way is C_A and C_B to Aircraft A and B, respectively, an efficient resolution does not give the right of way to A if $C_A < C_B$, and the total cost of the conflict is the minimum of C_A and C_B. To facilitate quantitative comparisons, we define *50% efficiency rate* as the outcome of a random allocation rule (the total cost of the conflict in this case is the mean of C_A and C_B) and assume a linear relationship between cost and efficiency. Thus *efficiency,* as defined here, can be thought of as normalized payoff (on an interval scale).

The most important problem is that different aircraft are likely to value the right of way differently ($C_A \neq C_B$). When the left-hand aircraft values the right of way more, granting it to the right-hand aircraft creates inefficiency. Assuming independence between costs and sides, the expected efficiency of a random assignment rule is 50%.

A second and less obvious inefficiency can occur because aircraft will be motivated to deviate from the direct route in order to win the right of way. Consider a situation in which pilots can select to approach a particular airport from one of two routes, left or right, and aircraft from the right have the right of way. Assume that (a) the flight time from the right route is more costly (in terms of time and/or money) by $R > 0$ units and (b) when two aircraft select the same side, a random event determines which one receives the right of way. Note that the right of way can be lost by either losing contention for a desired path or losing it to the aircraft on the right by rules of the road.

This problem is abstracted in Game 1 (Table 1, Part A). The rows indicate the route choice for Aircraft A, the columns for Aircraft B; cell entries reflect the costs of those choices. The notation (x, y) implies cost of x to Aircraft A and cost of y to Aircraft B. Thus the upper right hand cell of Game 1 shows the case in which A chooses the left route, B the right. Aircraft A incurs cost C_A for loss of right of way by rules of the road, and B incurs cost R for the lengthier route. The random events are abstracted with a gamble between two cost profiles. In the upper left cell, for example, the loser of the contention for the left route incurs a cost of C_i for loss of right of way, the winner a cost of 0. In the bottom right cell, the loser again incurs a cost of C_i for losing the right of way, but both incur a cost R for the lengthier right route. The numerical example in Part B of Table 1 demonstrates the implications of Game 1 with the cost parameters $R = 2$, $C_A = 15$, $C_B = 25$.

Rational agents. Expected utility maximizing (rational) agents are predicted to condition their behavior in Game 1 on the magnitude of R and C_i. (In the analysis of rational agents, we assume that the payoffs represent utilities. Thus risk aversion considerations do not apply. In addition, we focus on one-shot play of each game. The one-shot assumption seems reasonable, given the large number of aircraft and the small probability of repeated interaction.) Specifically, rational Aircraft i is expected to select the right route if $R < C_i/2$. Thus when $R < C_i/2$, the expected efficiency rate is below 50%. In that case, Game 1 is an example of the prisoner's dilemma game. Selection of the right route is rational to the individual, but it impairs the group payoff. The total expected cost of the conflict (for both aircraft together) is $2R + mean(C_A,C_B)$ (24 in the numerical example), whereas minimal conflict cost is obtained when both aircraft select the left route and the right of way is granted to the aircraft that values it more. The total conflict cost in this case is only $min(C_A,C_B)$ (15 in the numerical example). Because 100% efficiency involves a conflict cost of 15 and 50% involves a conflict cost of 20 (the mean of 15 and 25), the implied efficiency rate (from conflict cost of 24) is 10%.

Adaptive learning. To derive the expected behavior of adaptive agents, we ran computer simulations in which virtual agents that behave according to Erev and Barron's (2003) model repeatedly play the game presented in Table 1,

TABLE 1: Matrix Presentation of Game 1

	Aircraft B	
Aircraft A	Left	Right
	Part A[a]	
Left	(C_A, 0) with probability .5, (0, C_B) otherwise	(C_A, R)
Right	(R, C_B)	($R + C_A$, R) with probability .5, (R, $R + C_B$) otherwise
	Part B[b]	
Left	(15, 0) with probability .5. (0, 25) otherwise	(15, 2)
Right	(2, 25)	(17, 2) with probability .5, (2, 27) otherwise

[a]In Part A, the payoffs represent the cost of each aircraft as a function of the selected route (left or right) and of the random draw (on the diagonal). The notation (*x*, *y*) stands for payoff of *x* to A and payoff of *y* to B. [b]Part B has the parameters $R = 2$ and $C_A = 15$, $C_B = 25$.

Part B. We considered a pairwise interaction among 100 aircraft (50 with a cost of 15 and 50 with a cost of 25). On each trial all the aircraft were randomly paired. The predictions are presented in Figure 1 in terms of the efficiency in eight blocks of 100 trials each. The lower curve (labeled "right wins") summarizes the current case. It shows that the model predicts steady learning toward the problematic predictions of the rationality assumption (10% efficiency rate). Additional sensitivity analysis shows that this prediction is robust to the value of the cost parameters (under the constraint $R < C_i/2$).

POSSIBLE SOLUTIONS

In order to highlight the different properties of possible solutions to the right-of-way problem, the current analysis focuses on the case in which the costs (from losing the right of way) to Aircraft A and B (C_A and C_B) are drawn from the uniform (C_{low}, C_{high}) distribution with $0 < C_{low} < C_{high}$. Maximal (100%) efficiency in this context is obtained by granting the right of way to the aircraft that values it more. The expected cost of the conflict, in this case, is the expected value of the minimum of C_A and C_B. The known properties of the uniform distribution implies that this value is $E[min(C_A,C_B)] = C_{low} + (C_{high} - C_{low})/3$. A random rule (50% efficiency) is associated with an expected conflict cost of $(C_{low} + C_{high})/2$.

Free Negotiation With a Stated Cost Rule

In theory, pilots could negotiate the right of way. In a perfect world, free negotiation could lead to an efficient outcome in which the aircraft that values the right of way more would obtain it. However, under the assumption of rational or adaptive agents, free negotiation is unlikely to lead to efficient outcomes because pilots will be able to benefit from distorting their costs. To demonstrate this point, consider a rule that states that each aircraft has to state its costs (a number between C_{low} and C_{high}) and the aircraft that states the larger number wins the right of way. The costs in this simplified game, referred to as Game 2, are 0 for the winner and the actual cost (C_A or C_B) for the loser. The winner is determined with a random device in case of a tie.

Rational agents. The rational behavior in Game 2 is to state the maximal possible cost, C_{high}. This strategy is the dominant choice and the unique equilibrium of the game. As a result, the expected costs are identical to those of the random device.

Adaptive agents. The "stated cost" curve in Figure 1 presents the observed efficiency in a

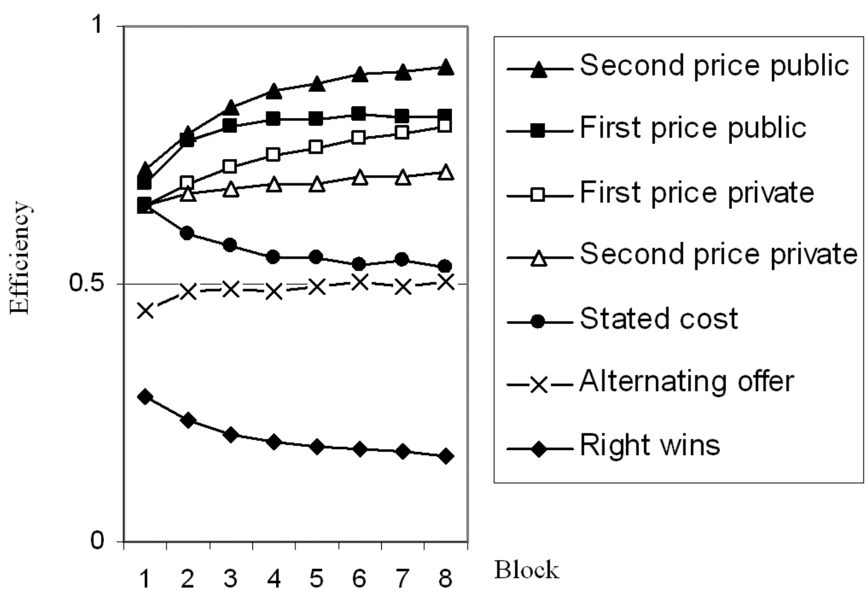

Figure 1. Efficiency rates in eight blocks of 100 trials.

Game 2 interaction between 100 Erev-Barron (EB; Erev & Barron, 2003) agents whose cost parameters were randomly drawn (at each trial) from the (5,35) interval. The simulation assumes that Agent i considers 13 strategies (stated costs). The 13 strategies include C_{low}, C_i, C_{high}, 5 strategies uniformly distributed between C_{low} and C_i, and 5 additional strategies uniformly distributed between C_i and C_{high} (sensitivity analysis demonstrates that the main results are insensitive to the assumed number of strategies). The results show relatively fast convergence of reported costs to the expected cost in equilibrium and of total costs to those expected using the random rule. That is, the simulated agents learn to inflate the reported costs. As a result the efficiency rate converges to 50% – the rate of the random rule.

Empirical observations suggest that this pessimistic prediction might be accurate. One example involves an academic department at one author's university in which financial support to graduate students was determined based on a subjective assessment of their relative performance. The students' advisers (who prefer that their students be supported) gave these assessments. Within 7 years, more than 50% of the students were ranked within the top 5th percentile.

An Alternating Offer Rule

Another difficulty that arises in negotiations relates to the fact that bargaining power may be related to patience rather than to the relevant costs. That is, the less patient aircraft may lose the right of way. To demonstrate this prediction, consider an environment in which the negotiation is restricted to a two-round alternating offer game. For simplicity, assume that the first offer can be "Please move" or "Go ahead, I'll move." The latter offer is always accepted, but the former can be rejected. In that case the winner of the right of way is determined by a random event, and the loser suffers an additional cost of $(C_i)(F_i)$ – that is, $(100)F_i$% of its cost parameter. Figure 2 presents this game in extensive form. Notice that this additional cost captures the idea of impatience that is known to affect bargaining power (see Rubinstein, 1982). We assume that the value of F_i is 0 for half the population (no impatience) and more than 1.5 to the other half (high impatience).

Rational agents. The backward induction equilibrium of Game 3 conditions behavior on the value of F_i. Patient agents ($F_i = 0$) are expected to reject the offer "Please move" as second movers and to make this offer as first movers. Impatient agents ($F_i > 1.5$), however, are expected to accept the offer to move as second mover and to volunteer to move as first movers. Under the assumption of independence of F_i and C_i, this prediction implies a 50% efficiency rate.

Adaptive agents. The "alternating offer" curve in Figure 1 presents the efficiency rate in Game 3

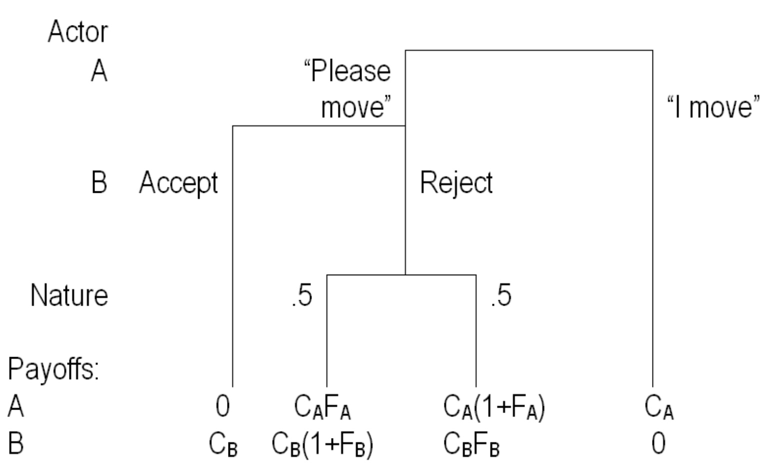

Figure 2. Extensive form presentation of Game 3. At the first stage the first mover (Aircraft A) selects an offer. Aircraft B can accept of reject the offer "Please move." In the case of rejection, a random event (nature) moves one of the aircraft. The payoffs represent expected costs.

when the players are EB adaptive agents. The F_i of the impatient aircraft was set to 2. The results show convergence to the problematic equilibrium efficiency.

Negotiation With Side Payments

In theory, these problems could be addressed by allowing side payments. That is, the high-cost aircraft can pay the low-price aircraft for the right of way. Unfortunately, however, this solution creates an incentive for low-cost aircraft to initiate conflicts (with a minimal course change at negligible cost) in order to "sell" the right of way. Thus, whereas this solution may be useful in resolving real conflicts, it might create artificial conflicts and, for that reason, impair efficiency.

The Sealed Bid Auction Solution

One solution to all the problems presented involves the use of sealed bid auctions. Under this solution, a decision maker at the airline level is responsible for keeping track of the expected cost of losing the right of way for each aircraft. These costs are updated on the central computer in the form of bids – the amount of "air money" that the airline is willing to pay for the right of way in the next conflict (if one occurs). When two aircraft approach each other in midair, the computer compares their bids (priority levels, costs) and gives the right of way to the aircraft with the higher bid. Two types of popular sealed bid auctions are considered here: "first price" and "second price" (see Vickery, 1961; Wilson, 1992). In a first price auction the winning airline pays its bid. In a second price auction the winner pays the second price (the bid of the loser). In both cases, the lower-bidding aircraft pays nothing, earns nothing, and loses the right of way. The amount paid by the winner goes to a third party.

Rational agents. First price auction has a unique equilibrium in which both agents bid a certain proportion of their true value. In the current setting (uniform costs), the equilibrium bidding for Aircraft A is $(C_A + C_{low})/2$. At this equilibrium the high-cost aircraft (m) always wins the right of way and pays $(C_m + C_{low})/2$. Under the assumption that the payment to be paid by the airlines will be put toward a common good (e.g., improving air safety), this solution ensures maximal efficiency. The expected conflict cost will be $E[min(C_A,C_B)] = C_{low} + (C_{high} - C_{low})/3$.

In a second price auction – also known as a Vickery (1961) auction – bidding the true value is a dominant strategy. This property ensures 100% efficiency, as in a first price auction: The right of way will be granted to the aircraft with the higher cost, with expected conflict cost of $C_{low} + (C_{high} - C_{low})/3$.

The fact that the payment will be made to a common good or a third party implies that the aircraft will not be motivated to initiate conflict (see the section titled Negotiation With Side Payments). In addition, both auction rules avoid the route selection problem discussed in the section titled Two Shortcomings of the Current Regulation. Because route selection will not affect the right of way, this decision should not be biased.

Adaptive agents. The top four curves in Figure 1 show the expected efficiency of the two auction rules under the assumption of EB adaptive agents with 13 strategies that were defined in the same manner as in the section titled Free Negotiation With a Stated Cost Rule, with Strategy 7 equaling C_i, which was independently sampled for each agent in each trial. We consider the 13-strategy case because it always includes the equilibrium strategy (it is $C_{low} + [C_i - C_{low}]/2$ in first price and C_i in second price). Sensitivity analysis shows that adding additional strategies does not change the results significantly. Agents were simulated assuming both public and private bidding. Public bidding implies that each player knows what the other bid is and, consequently, what outcomes could have been obtained had that player bid differently.

The results reveal learning toward efficient outcomes under both auction rules, with an interaction between the auction rule (first and second price) and the information available (public or private bids). When bidding is private there is an advantage to using the first price auction rule. This advantage is attributable to the perceived variance in outcome when using the second price rule; this perceived variance consequently impairs learning (Haruvy, Erev, & Sonsino, 2001). With public bidding the efficiency of both auction rules is increased, with a distinct advantage for second price (in Trials 701–800, the efficiency was 92% for second price and 82% for first price). The advantage of second price with public bidding is a result of the assumption that agents tend to follow a strategy of "best reply to recent outcomes." An agent playing this strategy bids according to the strategy that would have maximized returns in recently played rounds, had it been played. In first price, this assumption implies deviations from the efficient equilibrium (because the best reply can be to bid a little more than the other bidder). Deviations of this type do not occur in second price; in this setting the efficient equilibrium strategy dominates all other strategies and is always the best reply to recent outcomes.

CONCLUSIONS

Under the free flight program, a self-separation regulation that gives the right of way to the right-hand aircraft will be associated with two inefficiencies. First, in many cases the right of way would not be granted to the aircraft that values it most. Second, it might lead aircraft to deviate from the direct routes in order to win the right of way. The current analysis evaluates four possible solutions to this problem. The analysis suggests that some of the apparently reasonable negotiation-based solutions can create major problems: Rational or adaptive agents might distort their values (see Free Negotiation With a Stated Cost Rule), the aircraft with the higher final cost will lose the negotiation (see An Alternating Offer Rule), and the possibility of side payment will increase the likelihood of conflicts (see Negotiation With Side Payments).

To avoid these inefficiencies, we propose the use of auction rules (see The Sealed Bid Auction Solution). Under these rules aircraft (or airlines) would insert a priority bid to a central computer that would be used, in case of a conflict, to determine the right of way. Whereas first price and second price auctions are expected to be equally effective under the rationality assumption, second price auction with public bidding is more effective given adaptive agents. Although such competitive auctions seem novel in the context of air traffic control, they are consistent with recent FAA-sponsored attempts to improve the efficiency of airspace operations through collaborative decision making, fostered by increased information sharing. Auctions, or related approaches such as mediated bartering (Vossen & Ball, 2001), provide a means of sharing information about the value to individual players of certain outcomes.

Future Research

It is important to re-emphasize that the problems discussed here are only simplified examples of some of the difficulties expected to emerge in interactions between approaching aircraft. The current analysis ignores, for example, the role of air traffic controllers and specific systems such as TCAS. Nevertheless, the outstanding success of auction-related ideas in other mechanism design problems (e.g., Lucking-Reiley, 2000; McAfee & McMillan, 1996; Plott, 1997) suggests that the problems and solutions discussed here should be considered seriously. We believe it is possible that these ideas can be useful in more complex settings. For example, the priority bids proposed here can be used by air traffic controllers and/or be weighted in the prescription of new versions of TCAS.

Recent attempts to study the effect of flight regulations discovered some nontrivial strategic behaviors. For example, in one flight-simulator study (Lozito et al., 2000), an aircraft that was supposed to make an avoidance maneuver by the right-of-way rules specifically requested air traffic control to cancel those rules and make the other aircraft maneuver. Other studies (e.g., Endsley et al., 1997; Galster et al., 1998; Hilburn et al., 1997) have suggested that automated conflict detection might reduce the controller's awareness of the traffic situation. However, little effort was made to study the strategic interaction directly. We hope that the current analysis will facilitate deeper game theoretic, experimental, and empirical evaluation of the nontrivial economic incentives that can be created by new flight regulations.

APPENDIX

Erev and Barron's (2003) model is an extension of the model proposed by Erev, Bereby-Meyer, and Roth (1999). It can be summarized by the following assumptions.

RL1: Cognitive Strategies and a Two-Stage Choice Process

Two choices are made in each trial: The player first selects a set of strategies ("direct" or "cognitive"). At the second stage the player selects one of the strategies in the selected set. The direct set includes the alternatives explicitly presented by the game (e.g., *L* or *R* in Game 1). The cognitive set includes the following two strategies.

Hill climbing. This strategy (if used at trial t+1) prescribes a selection of the alternative with the highest weighted value in the set that includes the alternatives for which payoffs were observed in the last period and their neighbors.

If x_{bt} – the payoff of alternative b at trial t – is observed, then the weighted value of b is set to equal the mean of the new payoff and the previous record. That is, $WV_{b,t+1} = mean(WV_{bt}, x_{bt})$. If the payoff of b is not observed at t, then $WV_{b,t+1} = WV_{bt}$. The initial value is $WV_{b1} = A(1)$, the expected payoff from random choice. When a payoff from a neighbor of alternative b is observed before the first selection of b, b's record is updated based on the neighbor's payoff. Indifference is resolved with a random draw.

Loss avoidance. This strategy implies a selection of the alternative that consistently minimizes the probability of losses. If this rule implies indifference (e.g., because all alternatives always or never lead to losses), then the indifference is resolved by selecting one of the alternatives using the choice rule described in Equation 1.

RL2: Stochastic Choice Rule

The probability p_{kt} that option k in decision d (d = 1 involves the decision between the two sets, d = 2 is the choice among the strategies in the selected set) at time t is given by the choice rule used previously,

$$P_{kt} = \frac{\exp(\lambda_t q_{kt})}{\sum_{j \in J} \exp(\lambda_t q_{jt})}, \qquad (1)$$

in which λ_t is a payoff sensitivity term (defined in RL3) and q_{jt} is the propensity to select option j.

RL3: Reinforcement Updating

If option k is selected at trial t, the propensity to select it in trial t+1, $q_{k,t+1}$, is a weighted average of the propensity in t and the obtained payoff x_{kt}:

$$q_{k,t+1} = (1 - w_{dt})q_{kt} + (w_{dt})x_{kt}. \qquad (2)$$

The weight of the new reinforcement is $w_{dt} = 1/[N(1)/d + C_{kt}]$, in which $N(1)$ captures the "strength" of the initial value q_{k1} and C_{kt} is number of times k was selected in the first t trials. To capture the sensitivity of learning speed to the length of the experiment (T) and the number of alternatives (m), $N(1) = \min[\eta, T/(m - 1)]$, in which η is a free parameter. The initial propensity q_{j1} is assumed to equal $A(1)$. On each trial two propensities are updated, one for the chosen set and one for the chosen alternative or strategy. Propensities of options that were not selected are not updated.

RL4: Payoff Sensitivity

The payoff sensitivity at trial t is $\lambda_t = \lambda/S(t)$, in which λ is a payoff sensitivity parameter and $S(t)$ is a measure of observed payoff variability:

$$S(t + 1) = S(t)[1 - w'(t)] + AD(t)w'(t), \qquad (3)$$

in which $w'(t) = 1/[t + N(1)]$ and $AD(t)$ is the perceived relative payoff deviation in trial t. The exact value of $AD(t)$ depends on the available forgone payoff information,

$$AD(t) = \begin{cases} |R(t) - A(t)| \\ \textit{no forgone payoff information} \\ MN\big(|R(t) - A(t)|, |R(t) - R_k(t)|\big) \\ \textit{with forgone payoff information,} \end{cases} \qquad (4)$$

in which $A(t)$ is a measure of payoff average and act k is the act with the direct strategy that has the highest propensity among all the acts not selected at t. The average payoff measure, $A(t)$, is updated in the same way as the payoff variability term:

$$A(t + 1) = A(t)[1 - w'(t)] + R(t)w'(t). \qquad (5)$$

Altogether, the model has two learning parameters. Erev and Barron's (2003) estimation yields the values λ = 4.5 and η = 200. The prediction of the model for the current settings were derived by running computer simulations in which virtual agents that behave according to the model's assumptions play each of the games.

ACKNOWLEDGMENTS

This research was supported by a grant from NASA, the USA-Israel Binational Science Foundation, and the Technion Vice President Research Fund. We thank Ernan Haruvy, Al Roth, and the participants of seminars at Harvard, Columbia, New York University, University of Michigan, and University of Chicago for useful comments.

REFERENCES

Camerer, C. F., & Ho, T. (1999). Experience-weighted attraction in games. *Econometrica, 64,* 827–874.

Chen, Y., & Tang, F. (1998). Learning and incentive compatible mechanisms for public goods provision: An experimental study. *Journal of Political Economy, 106,* 633–662.

Cheung, Y. W., & Friedman, D. (1998). A comparison of learning and replicator dynamics using experimental data. *Journal of Economic Behavior and Organizations, 35,* 263–280.

Cooper, D., Garvin, S., & Kagel, J. (1997). Signalling and adaptive learning in an entry limit pricing game. *Rand Journal of Economics, 28,* 662–683.

Daniel, T. E., Seale, D. A., & Rapoport, A. (1998). Strategic play and adaptive learning in sealed bid bargaining mechanism. *Journal of Mathematical Psychology, 42,* 133–166.

Endsley, M. R., Mogford, R., Allendoefer, K. R., Snyder, M. D., & Stein, E. (1997). *Effect of free flight conditions on controller performance, workload, and situational awareness* (FAA Tech. Report). Atlantic City International Airport, NJ: U.S. Department of Transportation, Federal Aviation Administration.

Erev, I., & Barron, G. (2003). *On adaptation, maximization, and reinforcement learning among cognitive strategies.* Unpublished manuscript, Technion – Israel Institute of Technology, Haifa.

Erev, I., Bereby-Meyer, Y., & Roth, A. E. (1999). The effect of adding a constant to all payoffs: Experimental investigation, and implications for reinforcement learning models. *Journal of Economic Behavior and Organizations, 39,* 111–128.

Erev, I., & Roth, A. E. (1998). Predicting how people play games: Reinforcement learning in games with unique strategy equilibrium. *American Economic Review, 88,* 848–881.

Fleming, K., Lane, J., & Corker, K. (1999, December). Measuring controller reactions to free flight in a complex transition sector. *Air Traffic Control Quarterly* (n.p.).

Galster, S. M., Duley, J. A., Masalonis, A. J., & Parasurman, R. (1998). Effects of self-separation on controller conflict detection performance and workload in mature free flight. In M. W. Scerbo, W. Mark, & M. Mouloua, (Eds.), *Automation technology and human performance: Current research trends* (pp. 96–101). Mahwah, NJ: Erlbaum.

Haruvy, E., Erev, I., & Sonsino, D. (2001). The medium prizes paradox: Evidence from a simulated casino. *Journal of Risk and Uncertainty, 22,* 251–261.

Hilburn, B., Jorna, P. G., Byrne, E. A., & Parasuraman, R. (1997). The effect of adaptive air traffic control (ATC) decision aiding on controller mental workload. In M. Mouloua & J. Koonce (Eds.), *Human-automation interaction: Research and practice* (pp. 84–91). Mahwah, NJ: Erlbaum.

Hoffman, E., Zeghal, K., Cloerec, A., Grimaud, I., & Nicolaon, J. (2000). *Is limited delegation of separation assurance promising?* Paper presented at the Third USA/Europe Air Traffic Management R&D Seminar. Retrieved May 10, 2004, from http://www.eurocontrol.fr/projects/freer/archive/118_napoli.pdf

Livack, G., McDaniel, J., Battiste, V, & Johnson, W. (1999, April). *Human factors considerations in automatic dependent surveillance – Broadcast (ADS-B) flight operations.* Paper presented at the Fourth International Civil Aviation Organization Global Flight Safety and Human Factors Symposium, Santiago, Chile.

Lozito, S., McGann, A., Cashion, P., Dunbar, M., Mackintosh, M., Dulchinos, V., et al. (2000). *Free flight simulation: An initial examination of air-ground integration issues* (NASA Tech. Memorandum 2000-209605). Moffett Field, CA: NASA Ames Research Center.

Lucking-Reiley, D. (2000). Vickery auctions in practice: From 19th-century philately to 21st century e-commerce. *Journal of Economic Perspectives, 14,* 183–192.

McAfee, P. R., & McMillan, J. (1996). Analyzing the airwaves auction. *Journal of Economic Perspectives, 10,* 159–175.

National Civil Aviation Review Commission. (1997). *Avoiding aviation gridlock: A consensus for change.* Washington, DC: Author.

Plott, C. R. (1997). Laboratory experimental testbed: Application to PCS auction. *Journal of Economics and Management Strategy, 6,* 605–638.

Roth, A. E., & Erev, I. (1995). Learning in extensive-form games: Experimental data and simple dynamic models in intermediate term. *Games and Economic Behavior, 8,* 164–212.

Radio Technical Commission for Aeronautics. (1994). *Report of the RTCA Board of Directors' Select Committee on Free Flight.* Washington, DC: Author.

Rubinstein, A. (1982). Perfect equilibrium in bargaining model. *Econometrica, 50,* 97–109.

Sarin, R., & Vahid, F. (2001). Predicting how people play games: A simple dynamic model of choice. *Games and Economic Behavior, 34,* 104–122.

Stahl, D. O. (1999). Evidence based rule learning in symmetric normal-form games. *International Journal of Game Theory, 28,* 111–130.

Thorndike, E. L. (1898). Animal intelligence: An experimental study of the associative processes in animals. *Psychological Monographs, 2,* 1–109.

Vickery, W. (1961). Counterspeculations and competitive sealed traders. *Journal of Finance, 16,* 8–37.

Vossen, T., & Ball, M. (2001). *Optimization and mediated bartering models for ground delay programs.* Unpublished manuscript, Pennsylvania State University, University Park.

Wickens, C. D., Mavor, A. S., & McGee, J. P. (Eds.). (1997). *Flight to the future – Human factors in air traffic control.* Washington, DC: National Research Council/National Academy.

Wilson, R. (1992). Strategic analysis of auctions. In R. J. Aumann & S. Hart (Eds.), *Handbook of game theory* (Vol. 1, pp. 227–279). Amsterdam: North-Holland.

Ido Erev received his Ph.D. in psychology from the University of North Carolina in 1990. He is the graduate studies coordinator for the Faculty of Industrial Engineering and Management at Technion – Israel Institute of Technology.

Greg Barron received his Ph.D. in behavioral science from Technion – Israel Institute of Technology in 2003. He is the Computer Lab for Experimental Research research fellow at Harvard Business School.

Roger W. Remington received his Ph.D. in psychology from the University of Oregon in 1978. He is group leader for cognition in the Human Factors Research and Technology Division at NASA-Ames Research Center.

Date received: January 18, 2002
Date accepted: December 29, 2003