Game-Theoretic Methods for Analysis of Tactical Decision-Making Using Agent-Based Combat Simulations

Dr. Yevgeniy Vorobeychik

University of Pennsylvania, yev@seas.upenn.edu

Dr. Isaac R. Porche, III

RAND Corporation, porche@rand.org

APPLICATION AREAS: Battle Management/ Command & Control; ISR & Intelligence Analysis

OR METHODOLOGIES: Linear Programming; Decision Analysis; Simulation

ABSTRACT

In this paper we provide a motivation and methods for using game-theoretic analysis for combat simulations. Application to three previously studied scenarios using the MANA (Map-Aware Non-Uniform Automata) simulation tool demonstrates the efficacy of our approach, and we provide a concrete example of the differences between the game-theoretic and the more standard decision-theoretic approaches. We believe that the game-theoretic simulation analysis framework can enhance the understanding of tactical combat landscape, especially in cases where intelligence efforts are ineffective or provide unreliable information, because in those cases game-theoretic analysis may help to corroborate or challenge intelligence-limited beliefs by assessing the rationality of the opponent's apparent strategy. We suggest that game-theoretic simulation analysis could often best be used during the low-resolution, exploratory-analysis phase of a multiresolution analysis.

INTRODUCTION

The increasing complexity of modern warfare imposes an unprecedented level of demand upon the efficacy of Command, Control, Communications, Computers, Intelligence, Surveillance, and Reconnaissance (C4ISR). Additionally, as the networking tools are becoming more advanced, the techniques to process the information and make effective decisions based on the results are becoming ever more precious. The use of computing power to aid the decision-making process must thus become an important component of military decision-making.

The advent of computer-based battlefield simulation models added an important tool to aid military analysis, and battlefield simulations have now become widespread in that capacity (for example, Horne, 1997; Ipekci, 2002; Porche et al., 2005, 2007; Porche and Wilson, 2006). However, few battlefield simulation analysis techniques incorporate game-theoretic methods. Granted, the questions that can practically be studied using simulations are inherently limited, and thus it is often reasonable to provide only high-level information and guidelines based on simulation results. However, as simulations increase in sophistication over time and as military decision-makers are faced with an increasing number of choices and increasing complexity surrounding them, we anticipate that computational decision methods based on simulations will grow in importance and effectiveness.

This work is a step towards a methodology of simulation data analysis that provides concrete probabilistic guidelines regarding best strategies for military operations. Our analysis takes into account a critical component that has generally been missing from combat simulation analysis in the past: the effect of rational decision-making of adversaries on each other's optimal strategies.[A] We do so by computing a Nash equilibrium (Nash, 1951), which is defined to be a set of competing strategies such that no side has an incentive to change its strategy unilaterally, based on simulation data. In contrast with the more traditional approaches that hold the tactics of opponent forces (e.g., Red) fixed, we allow opponents to select a strategy that is a rational response to the decisions by a friendly force (e.g., Blue). Moreover, our analysis incorporates uncertainty about numerous random (unknown) elements of the battlefield scenario. Analysis in the same spirit was done by Yiu et al. (2002) to investigate strategies to manage civil violence. However, the use of game theory in their study was relatively rudimentary and context-specific, with no attention paid to the stochastic nature of simulation outcomes. Several other game-theoretic analyses of combat scenarios have been undertaken in literature (for example, Haywood, 1951, Hillestad, 1986, Lachman and Hillestad, 1986, Hamilton and Mesic, 2004), although none of these actually undertake a study of simulation outcomes. Very much in the spirit of our work, Reeves et al. (2005) and Wellman et al. (2008) present a set of methods for game-theoretic analysis of simulations in the context of simultaneous ascending auctions. Walsh et al. (2002) and Schvartzman and Wellman (2009) have used similar analysis techniques in the context of continuous double-auctions.

We chose MANA (Galligan et al., 2004) as our combat simulation tool. MANA-based models such as ours introduce some artificialities that limit the ultimate validity of our analysis (see Appendix), but they are suitable to illustrate the larger points that we wish to make. Furthermore, many of the limitations of MANA can be alleviated via a multi-resolution modeling framework, in which strategies obtained based on a low-resolution model are subsequently reevaluated using a high-resolution model. The main purpose of this work is, thus, not to prescribe optimal strategies on the particular battlefield contexts we explored, but to advertise a set of tools that can be used to produce such results when all the proper adjustments and qualifications have been made. In what follows, we present a rigorous simulation-based game-theoretic analysis of three MANA combat scenarios. For each, we derive equilibrium outcomes and strategies for Blue and Red in the particular restricted strategy spaces. Additionally, we evaluate the sensitivity of our results to simulation noise. We further complement the technical results with high-level intuition. For two of the three scenarios, we considered both a one-stage model, in which the players make their decisions simultaneously, and a two-stage model, in which Blue first chooses one of its decision variables, and then both Blue and Red (upon observing the choice of Blue) make the remaining strategic choices simultaneously.

ANALYSIS OF MANA SIMULATIONS

Parameter Choices

The first step in applying empirical game-theoretic methods to simulation data is to specify the strategies of the players. Within MANA we can specify parameters that define the behavior of Blue and Red agents on the simulated battlefield. Two examples of such parameters are (i) tendency to approach friendly or enemy units and (ii) distance at which another entity is recognized.

Suppose we can adjust B parameters for Blue agents and R parameters for Red agents. We define a strategy of Blue to be a particular setting of its B parameters, and similarly, the strategy of Red will be a setting of its R parameters. The respective strategy sets of Blue and Red will then be all possible settings of their parameters. A Nash equilibrium is computed for a game in terms of strategies, ^{B} and we map it to the corresponding parameter settings as follows. Let $\Pr\{B_i = b_i\}$ denote the probability of Blue setting its ith parameter to a value b_i in equilibrium. We can show that this probability is the sum of the probabilities of strategies in which $B_i = b_i$:

$$\begin{aligned}\Pr\{B_i = b_i\} &= E\,\Pr\{B_i = b_i | s_b = \alpha\} \\ &= \sum_{\alpha \in S_b} \Pr\{B_i = b_i | s_b = \alpha\} \Pr\{s_b = \alpha\},\end{aligned}$$

where

$$\Pr\{B_i = b_i | s_b = \alpha\}$$

is 1 when $B_i = b_i$ in strategy α and 0 otherwise.

Since parameters in MANA correspond to simple low-level behavior or capabilities of battlefield agents, the analyst must carefully select parameters that have meaningful interpretations within an actual battlefield. For the scenarios we study here, we chose to model the following behavioral properties of agents: *aggressiveness* (by setting an agent's tendency to move toward alive enemies), *cohesion* (by setting the tendency to move toward alive friendly forces), and *sensing* (by setting the range at which enemy, friendly, or neutral forces are detected and identified).

Aggressiveness models an agent's (soldier's) reactive behavior towards encountered enemy units. We can view aggressiveness as a reflection of orders to engage enemy units upon encounter. A non-aggressive agent, on the other hand, may be following orders to avoid enemy in order to achieve its goal as covertly as possible. Aggressiveness may also reflect a type of training that soldiers may receive. Additionally, this parameter can model intangible character traits such as patriotism or fear. In the reported experiments, we considered three settings of aggressiveness. We refer to the first setting as *evasive* (MANA enemy encounter setting of −100), alluding to the tendency of agents to avoid confrontation with enemy agents. The second setting is *passive* (MANA enemy encounter setting of 0), which suggests that agents do not respond to encounter with unfriendly forces, moving neither toward nor away from them. The third setting is *aggressive* (MANA enemy encounter setting of

100), which imbues agents to approach the enemy forces upon encounter.

Cohesion is the proclivity of an agent to remain close to other friendly units. This may model the training that instructs soldiers to stick together to avoid being singled out by small mobile enemy units. (We do not necessarily mean that soldiers would be immediately next to one another, but simply in reasonable proximity.) On the other hand, lack of cohesion may correspond to commander instructions to engage in diversionary tactics or split up a search party into small but highly mobile units. Additionally, cohesion may reflect intangible traits such as camaraderie and morale amongst soldiers. We differentiate between three settings of cohesion. When the corresponding MANA parameter is set to -100, we call the agents *isolated*, referring to their tendency to stay away from detected friends. *Indifferent* agents do not react in any way to detected friendly forces. Finally, we call a force *clustered* when the agents tend to remain close to other friendly agents.

Sensing is a capability of detecting and identifying other agents. This capability may map to the appropriate sensing equipment available to solders on the battlefield, or to individual characteristics such as eyesight. Additionally, it may reflect training that helps identify enemy soldiers who attempt to blend in with civilians. The lowest setting of sensing that was used in our experiments was 10, which we simply refer to as *low*. A *medium* setting of sensing was 30, whereas the final setting we allowed was 50, termed *high*.

Setup

Opposing forces clashing on a battlefield can be described as zero-sum games (von Neumann and Morgenstern, 1953) in which one side's gain is the other side's loss. While the situation is more complicated in practice, this aspect is an instructive focus, as it is generally quite significant.[C]

It is common for studies and analyses of combat simulations to use Expected Loss Ratio and/or probability of success as measures of effectiveness (MOEs) of the Blue force (e.g., Lauren, 2001; Porche et al., 2005, 2007; Porche and Wilson, 2006); MOEs are designed to evaluate the accomplishment of mission objectives and achievement of desired effects (JCDRP, 2004). To use only the former means to ignore entirely the fact that missions are generally focused around specific goals and are typically considered unsuccessful if the goal is not reached. On the other hand, considering only the probability of success ignores the importance of achieving the goal with a minimum casualty rate. Thus, in this study, we combine the two metrics and introduce a parameter, w, that determines their importance vis-à-vis each other.

Since our study is game-theoretic in nature, we need to define utilities (real-valued functions corresponding to underlying preferences) for both players. However, the fact that our model is a zero-sum game imposes a restriction that utilities of players sum to zero, thereby eliminating the need to specify both, since one implies the other. In this work, we will generally refer only to the utility of Blue, leaving the utility of Red implicit. Formally, we define the utility function of Blue to be $u_b = \mathrm{LR} + w\mathbf{I}\{\text{Blue reaches goal}\}$ and the utility of Red to be $u_r = -u_b$. LR refers to loss ratio (adjusted to avoid a singularity when there are no Blue casualties): $\mathrm{LR} = \frac{\text{Red Casualties} + 1}{\text{Blue Casualties} + 1}$. $\mathbf{I}\{\text{Blue reaches goal}\}$ is 1 when Blue reaches its goal and 0 otherwise. Finally, w is the weight parameter that indicates the relative importance to Blue of reaching the goal as compared to the Loss Ratio. A large value of w indicates that reaching the goal is very important to the Blue force, even at the expense of a high casualty rate. A value of $w = 0$ indicates that reaching the goal is irrelevant to Blue, and small values suggest that Blue places a low value on the goal if many casualties can be sustained in the process of reaching it. Since there is considerable uncertainty about the outcomes (i.e., Loss Ratio and probability of success) even when all the parameters played by both Blue and Red are specified, the task of both sides would be to maximize their respective *expected* utility functions. Thus, for Blue the problem would take the form of $\max_{s_b \in S_b} (\mathbf{E}[\mathrm{LR}(s_b, s_r)] + w\mathbf{p}\{\text{Blue reaches goal} \mid (s_b, s_r)\})$, where s_b is the strategy of Blue, s_r is the strategy of Red (fixed in this optimization problem), and S_b is the strategy set of Blue. Red, on the other hand, would be minimizing this same quantity (for a fixed s_b) over its own strategy set S_r. Note that both the Loss Ratio and Blue probability of success are functions of the strategies of both players.

After we appropriately restrict the strategy space of the players, we run a fixed number of simulations for each strategy profile and average the results. Upon thereby acquiring an empirical payoff matrix, the first order of business is to find an empirical Nash equilibrium. When the game has few strategies played by both sides, it can be analyzed analytically, and the equilibrium can be found as a closed-form function of w. The analysis can then proceed to discuss the meaning of the equilibrium strategies, as well as the meaning of any *change* in equilibrium strategies or utility as w changes. The ''Future'' scenario below was small enough to allow such detailed analysis.

When the number of strategies is relatively large, however, we need numerical tools to analyze the resulting empirical games. Since an equilibrium solution of a zero-sum game can be formulated as a linear program, we used ILOG CPLEX (IBM, 2006) as the equilibrium solution tool. In order to analyze the relationship between equilibrium outcomes and strategies (or the corresponding parameter settings) in terms of w, we ran ILOG CPLEX using the payoff matrices for a discrete set of w.

''Future'' Scenario

This scenario pitted a small Blue force, consisting of two squads, against a Red force, which consisted of 100 dismounted fighters. (The term *squad* is used conceptually in MANA to represent any number of soldiers sharing the same characteristics.) A squad of six Blue infantry armed with direct fire weapons is moving from a position in the northwest to the southeast corner (see Figures 1a). They have the support of an indirect precision weapon with enough range to cover the entire city and a communication link that allows the Blue infantry squad to pass target information back to the indirect fire squad. The Blue force structure is illustrated in Figures 1b. Red forces were randomly placed throughout the map.

For our purposes, movement speed was left at one cell per simulation time step. We calibrate this to real-world time by considering that average human motion from walking to running is approximately 1-6 m/s. We assume that each cell is approximately five meters in length, that the agents move one cell per time step, and that

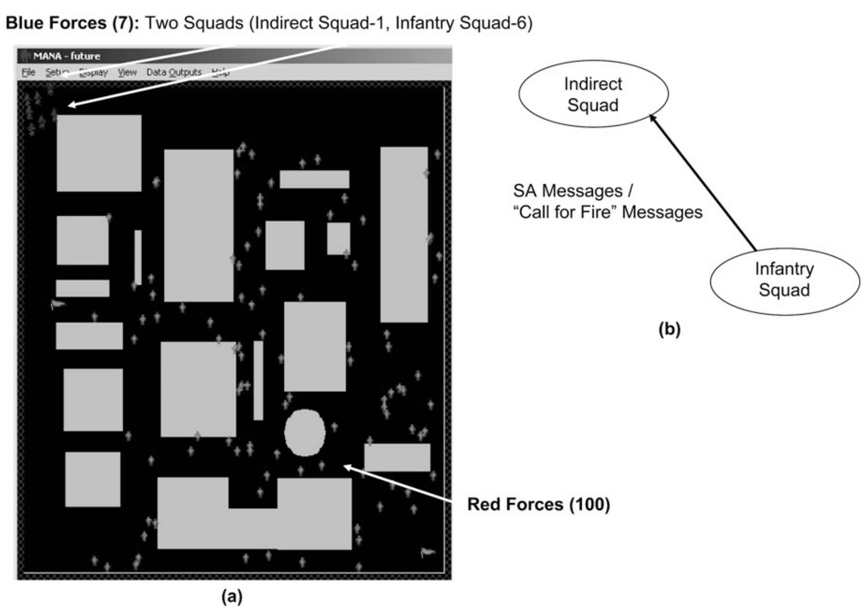

Figures 1a & b. MANA ''Future'' scenario.

a time step is one second in duration. Since our scenario environment was 200 $\times$ 200 cells, the approximate size of this fictitious urban environment is 1 km^2.

We defined the strategies of Blue and Red as their respective setting of the cohesion parameter. We then simulated 2,000 runs for each strategy profile (nine in total), with the averages of the Blue utility function values comprising the resulting payoff matrix of Blue. The resulting empirical payoff matrix is shown in Table 1. Estimated payoffs based on MANA ''Future'' scenario, with rows corresponding to different strategies of Blue and columns corresponding to Red's strategies. In each cell we record the average utility value of Blue when Blue and Red play their corresponding strategies. For example, when Blue and Red jointly play -100, the utility of Blue is $7.17 + 0.14w$. Observe that the utility of Blue is specified as a function of w. Consequently, we can incorporate this parameter in our analysis.

First, we determine the (empirical) optimal strategy of Blue when Red's cohesion parameter is its default setting of 0 (indifferent). In this case, Blue's best response is to be isolated. We can interpret this result as follows: When Red is moderately spread out (Red agents are indifferent between staying close or away from other Red agents), Blue will gain by splitting its mission up and allowing some proportion of forces to distract Red troops (Red will generally try to attack detected Blue troops), while the rest will be more likely to slip by unnoticed toward the goal, thereby maximizing probability of success and minimizing casualties. This result is independent of w, although it is important to recognize that even when w is zero, it does factor indirectly into the casualty rate, since the simulation ends as soon as Blue reaches the goal, thereby sparing Blue any additional casualties. (The analysis is somewhat more complex, since Blue is better armed but outnumbered, and consequently, the effect could be in either direction.)

Now we will allow Red to change its strategy in anticipation of Blue's decision and payoffs. We can readily observe from Table 1 that both the default setting of 0 (indifferent) and the setting of 100 (clustered) are strictly dominated for Red, that is, Red will not choose these strategies for any possible strategy of Blue. The effect of this reasoning is to cause Red's 100 troops to disperse (the ''isolated'' strategy in Table 1) throughout the battlefield, and they will likely remain relatively dispersed even when Blue ''diversionary'' agents are encountered. Thus, the strategy of Blue that calls for a high level of dispersion (as discussed above) will no longer be very effective. Instead, the best response of Blue in this case is to form a tight (clustered) group, taking advantage of its superior weaponry.

It is important to note here that the high-level conclusion we just made is not obvious prior to performing the empirical evaluation of the strategic landscape. The reasoning that we used to justify Blue being isolated when Red strategy is fixed at default can still in principle apply when Red is allowed to switch to -100 (isolated). It is only through the analysis of data that we were able to conclude that the benefit of forming a tight fighting unit was greatest.

In order to determine the significance of this conclusion, we assumed that the entries in the payoff matrix in Table 1 come from a Normal distribution (since they are means, they are asymptotically normally distributed, and with 2,000 samples per profile, the assumption is likely to be quite reasonable) and used the probabilistic bounds from Vorobeychik (2009). The results of this analysis (for w varying between

Table 1. Estimated payoffs based on MANA ''Future'' scenario

		Red		
		isolated	indifferent	clustered
Blue	isolated	$7.17 + 0.14w$	$12.48 + 0.60w$	$18.54 + 0.89w$
	indifferent	$7.09 + 0.14w$	$11.57 + 0.53w$	$15.89 + 0.80w$
	clustered	$8.08 + 0.15w$	$11.83 + 0.39w$	$15.16 + 0.68w$

0 and 100) suggest that we have no problem with statistical significance of our results in this scenario: for almost all settings of w the empirical equilibrium is an actual equilibrium with probability above 95%.

''Mosque'' Scenario: Direct Fire Weapons Only

This scenario involves a small group of Red leaders attempting to enter a mosque (hence the name) under the protection of Red ''leader protection'' troops (Figure 2). The Blue forces (four 6-man infantry squads) armed only with direct fire weapons try to prevent Red from achieving their goal. Blue forces have the support of UAVs that have stronger sensing capabilities. To keep Blue from firing on targets inside the mosque, a special terrain type was added that gave the Red forces 100% concealment while inside. There are two squads of Red forces, 6 leaders, and 30 security forces. The security forces are equipped with direct fire weapons, but with lower hit rate than Blue. They succeed (Blue fails) when the Red leaders can move unimpeded to the mosque.

We engaged in a somewhat more extensive analysis of this scenario, attempting to use empirical game theory to complement the analysis of Porche and Wilson (2006) and Porche et al. (2005, 2007) of the effect that sensing capability of Blue infantry troops (henceforth, Blue sensing) has on the effectiveness of Blue (in terms of the Blue utility function as defined above). We restricted the domain of Blue sensing to 10 (low), 30 (medium), and 50 (high) and selected two additional parameters for Blue and Red: aggressiveness and cohesion. The domain of aggressiveness and cohesion for both players was restricted as we had already described. The simulations were run 600 times for each strategy profile, and the average utility values were recorded as the estimates of the corresponding expected utilities. (The lower number of replications as compared to the ''future'' scenario is mainly due to the considerably increased simulation time.)

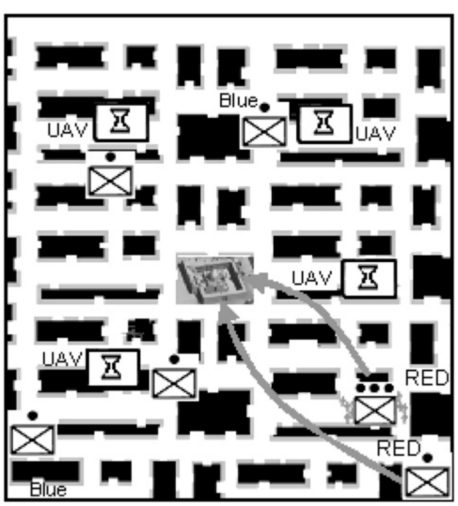

Figure 2. MANA ''Mosque'' Scenario. (Black boxes represent obstacles [e.g., buildings].)

Our choice of complementary parameters was driven primarily by the observation that sensing capability only comes into play in MANA when agents react to the resulting Situational Awareness (SA) map, and this effect is manifested in their tendencies to move toward or away from other agents detected. Thus, we expect the optimal choice of the aggressiveness and cohesion parameters of both players to change when sensing parameter value changes.

We study two alternative models of sensing. In one, Blue chooses its sensing parameter simultaneously with all the others and simultaneously with the strategy choice of Red. We call this the simultaneous choice model. In the second sequential choice model, we model the strategic interaction as a two-stage game (Figure 3). In stage one, Blue first selects its sensing capability. In stage two, Red observes the sensing choice of Blue, and both Blue and Red set the values for the remaining parameters simultaneously.

The most important difference between the two models is that the observation of strategic choice of Blue in stage one conditions the strategy choice of both players in stage two. The former model may correspond to the training of Blue troops (in secret), with Red unable to gather intelligence regarding the specifics of the training. The latter model is more appropriate when sensing capability models the actual sensing devices used by the troops that are well known. It may also (more generally) model effective intelligence efforts by the Red.

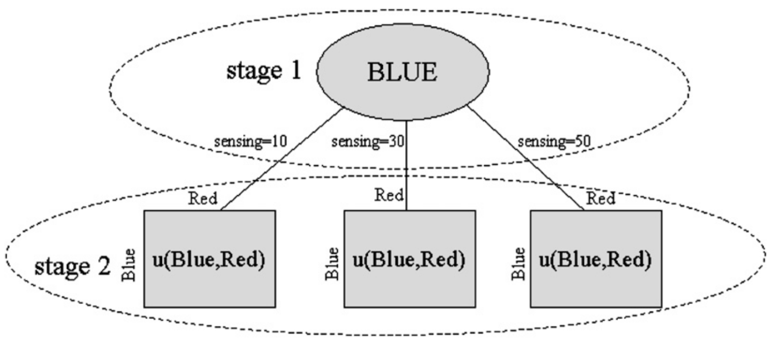

Figure 3. Sequential choice model.

In the analysis that follows, we present all of our results in terms of Nash equilibrium outcomes (loss ratio [LR] and probability of success in equilibrium) and strategies (parameter settings). As a part of the analysis, we varied the relative importance of mission success to Blue (variable w), and present our results as functions of this variable.

Simultaneous Choice Model. Recall that in the simultaneous choice model, Blue and Red are seen to select values for all parameters over which they have control simultaneously. Thus, our equilibria may (and will in general) involve either or both sides "mixing," that is, playing several values for each parameter with varying probabilities to avoid exposure to effective opponent response.

Figure 4a shows the plot of LR, and Figure 4b shows the probability of success of Blue in equilibrium as functions of w. We do not necessarily expect the relationships to be monotone, since a number of factors, including the opponent's (Red's) change in strategy, contribute to the resulting equilibrium outcome. Indeed, there are some intriguing non-monotonicities that we can observe. Consider, for example, the outcomes for $w = 6.5$ and $w = 7$ (circled in Figures 4a & b). In this instance, even though the importance of reaching the goal to Blue increases, the probability of success actually decreases slightly, although LR increases considerably. A result like this adds value to an additional strategy that Blue can entertain that deceives the adversary into believing that success is more important to Blue than it actually

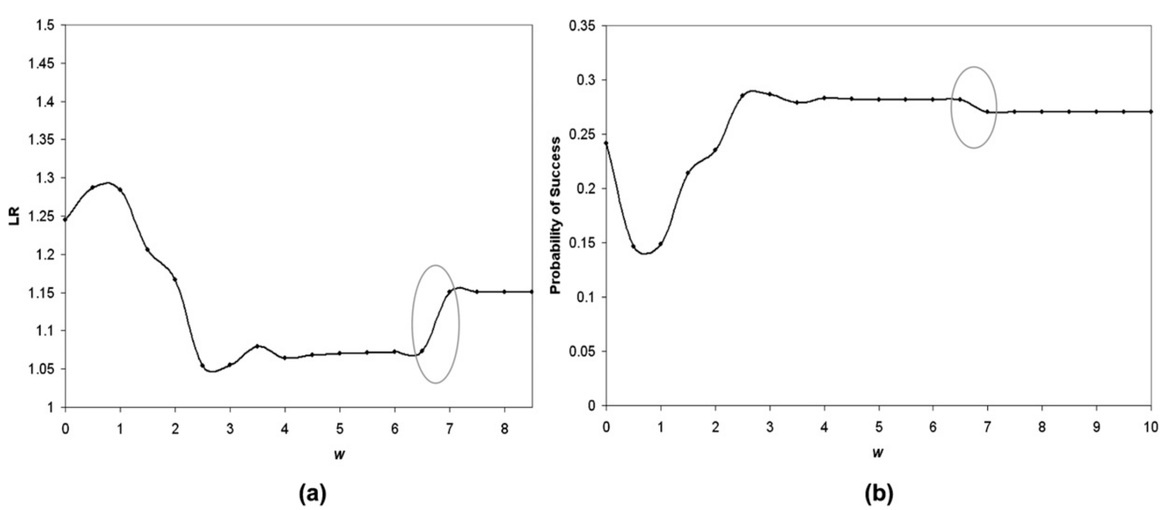

Figure 4. Loss Ratio (a) and Probability of success of Blue (b) as functions of w.

is in order to get the enhanced LR at little loss in probability of success. A similar non-monotonicity can be observed for $w = 0$ and $w = 0.5$. Here a deception strategy that leads the adversary to overvalue the importance of success to Blue may be in order. The availability of such deception strategies, of course, may dramatically change the nature of strategic interactions between the agents and we would need a careful study of an appropriate model of these to make concrete conclusions.

We can additionally observe that when w is high (above 7) its precise setting no longer affects the outcomes. We will see that equilibrium strategies stabilize in this region as well, although they tend to be highly dependent on the value of w when it is below 7.

As in the ''Future'' scenario, we would like to evaluate the statistical significance of our results, and to do this we apply the results from Vorobeychik (2009) for the range of w between 0 and 10. For each setting of w, we compute an equilibrium of the resulting empirical game and evaluate the probability with which it constitutes an actual equilibrium of the game (i.e., when actual expected payoffs are available). Figure 5a presents the resulting plot, and we can see that probabilities tend to be relatively low. It is somewhat expected that they will decrease as w increases, since the value of w provides a squared contribution to the variance of the empirical results. However, this relationship need not be monotone, since equilibrium strategies and outcomes do not necessarily vary nicely with w.

Although the probability that our empirical equilibria constitute actual equilibrium strategies and outcomes is relatively low, we have another way of measuring the quality of our approximation: ε-bound. We define ε-bound as the smallest ε such that the maximum benefit that any player can get by deviating to an alternative strategy is below ε. If this metric is small, we would expect that the outcomes and strategies based on empirical data are not very different from actual equilibrium results. (This need not hold in general, but we conjecture that it is true in typical games.) In order to normalize this metric, we divide it by the equilibrium utility of Blue, thereby producing the results seen in Figure 5b.

The plot in Figure 5b is more promising than Figure 5a, as it suggests that although our analysis deals with strategies and outcomes that are unlikely to be *exact* equilibria, they are likely to be *approximate* equilibria, with maximum benefit to deviation between 6% and 10% of empirical equilibrium utility with 95% confidence. Nevertheless, the high degree of uncertainty in the results gives us pause in making far-reaching conclusions based on these. Instead, we would like to suggest that our results provide an

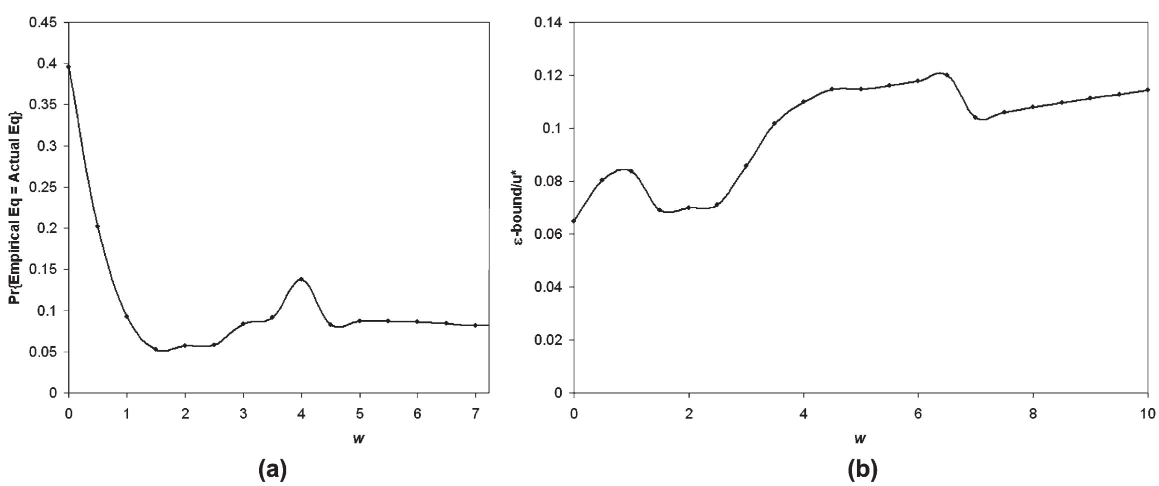

Figure 5. Probability that the equilibrium of the empirical game is an equilibrium of the actual game (a) and maximum benefit for deviation to any player with 95% confidence (b).

important insight into the nature of the strategic interactions between the adversaries in this scenario, and further explorations should be made into these.

For the analysis of equilibrium strategies, we assume that our equilibrium outcomes based on the empirical data represent the *actual* outcomes, and we ignore the issue of uncertainty about these we had just now addressed. Our main reason for doing so is that, given the data, we want to prescribe the best choices to the decision-maker. If we had more data, our conclusions would be more significant, but the analysis would proceed in much the same fashion.

Figure 6a shows the equilibrium probability with which all settings of the sensing parameter will be played by Blue. Our results here are essentially bi-modal: When w is low, medium sensing dominates, whereas for high values of w, Blue will select the low setting of sensing. It is most perplexing that Blue will never choose the highest possible sensing capability in spite of the fact that this comes at no *direct* cost (i.e., our model does not impose a cost on Blue for selecting a better sensing capability, even though realistically there would be some cost, either through retraining or due to purchasing the necessary equipment). Indeed, the lowest possible capability is selected for all values of w between 3 and 10. This result strongly contradicts our intuition, which suggests that our outcomes should never be worse when greater capabilities are available to us no matter what the opponent decides to do — this may not be the case in general games, but it is so in zero-sum games. Our intuition is based, among other things, upon the assumption that *everything else is equal*, which, it turns out, is not the case here. Particularly, we observe that actions available to players given equivalent information are different for different settings of sensing. We will discuss this issue in some detail later.

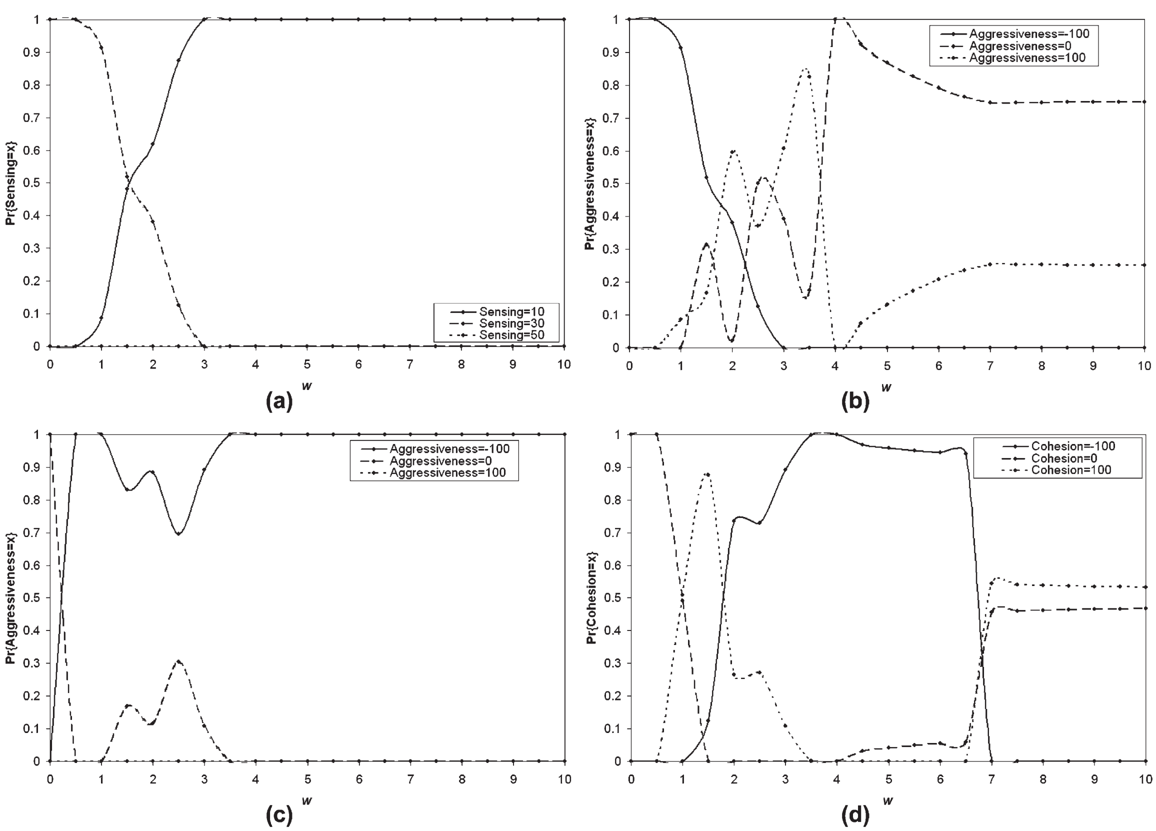

Figure 6. Sensing (a) and aggressiveness (b) of Blue. Aggressiveness (c) and cohesion (d) of Red.

Figure 6b displays the equilibrium probability distribution over different settings of aggressiveness of the Blue force, which is highly dependent on the value of w. When probability of success is unimportant, Blue will be evasive, avoiding encounters with Red troops. But even when the probability of success has importance, aggressiveness is relatively low (Blue is only occasionally aggressive). This result is likely due to the distracting effect of aggressiveness, which leads Blue troops to overcommit to the early inflows of Red forces, distracting them from the ultimate goal of keeping the leaders, who are generally trailing behind, from getting into the mosque. On the other hand, Blue must be aggressive with some probability, since it would otherwise not engage Red leaders at all.

Cohesion of Blue is high for all values of w in the range (thus, we do not display it in a plot). Thus, Blue will form tight squads in equilibrium to prepare for likely encounters with well-equipped Red troops.

Figures 6c and 6d display the equilibrium behavior of Red forces. Cohesion of Red is highly affected by w. When $w = 0$, Red will ignore friendly agents, presumably to avoid "putting all eggs in one basket" in instances of encounters with Blue. When w is high, Red forces will tend to make relatively close formations with roughly 50% probability (no reaction to SA the rest of the time), possibly to increase probability of repeated encounters that provide a sufficiently long distraction to allow Red leaders to "sneak by" into the mosque. Aggressiveness appears to settle for $w > 3.5$, with Red becoming evasive. Furthermore, for no value of w does Red behave aggressively. This result is quite intuitive since Red will minimize its losses by avoiding confrontation and, as it turns out, maximize the chances of Red leaders entering the mosque unharmed.

Let us now try to put everything together. When probability of success is unimportant ($w = 0$), Blue disperses its troops and avoids encounters with the well-equipped Red forces, thereby limiting its casualties. Medium level of sensing allows Blue a sufficient detection capability to attain the above objectives, but a greater capability (sensing=50) may actually work against Blue agents, who may become overly impulsive, responding to many stimuli at once and as a result possibly moving to undesirable locations with too many Red troops. In response to Blue, Red will avoid reacting to sensing readings entirely. A deviation to high cohesion may reduce its probability of encountering Blue troops, whereas low cohesion may result in Blue forming tight units and eliminating individual Red soldiers. Similarly, high aggressiveness may warrant a Blue response of forming very clustered units that increase their survivability in encounters against small groups of Red agents, whereas a high level of evasiveness may force Red into undesirable encounters with undetected Blue troops.

When w is 10, Blue will reduce its equilibrium sensing capability to 10 (lowest setting), remain dispersed, and will generally ignore Red troops when those are detected. In response, Red will be evasive and will form tight units about 50% of the time, ignoring them the other 50% of the time. Consider deviations in which Blue has higher sensing. By the same rationale as above, high sensing may cause Blue agents to be overly impulsive. The fact that Red forces become more likely to reach their goal with increased sensing levels of Blue suggests that reacting to other agents on the map provides too much distraction for Blue, which should instead essentially remain fixed around the possible entrance areas into the mosque in order to stand guard against Red leaders. Occasionally, Blue can then afford to be aggressive, since low sensing capability suggests that Red agents detected are close to entering the mosque and may need to be stopped. Since achieving the goal is now of prime importance to Red, they will avoid Blue at all costs, trying to move around them into the mosque. Much of the time, Red will form close units that are capable of drawn-out battles if approached by Blue, allowing Red leaders to enter the mosque during the window of opportunity provided by such battles.

Sequential Choice Model. The sequential choice model involves Blue first making its strategic choice regarding the setting of its sensing parameter, which is observed by Red. Thereafter, both Red and Blue select their remaining parameter settings and play them "simultaneously." (The play is simultaneous as long as no player can obtain any information about the strategies of other players.)

We model this particular mode of interaction as a two-stage game, in which Blue moves first by setting its kth parameter to b_k, and Blue and Red move together thereafter by playing a game in which Blue now has a restricted strategy set, selecting values for all parameters other than parameter k. Since b_k is observed, the strategic choices of both Blue and Red are now functions of b_k, as the value of b_k will generally change the mutually optimal choices for the players. In formal notation, we define the set of choices that Blue has for b_k in the first stage to be B_k. In the second stage, Blue can no longer choose a setting for its kth parameter, but can still set others, whereas the strategy set of Red remains S_r. Since the actual strategies of players will be conditioned on b_k, we define the set of Nash equilibria of the game <u>*in the second stage*</u> as $s^*(b_k)$ and the corresponding utility of Blue in equilibrium to be $u^*(b_k)$. (We need not specify which equilibrium strategy will be played because the equilibrium value is unique in a zero-sum game.) It is not difficult to show that if $b_k^* = \arg\max_{b_k \in B_k} u^*(b_k)$, then $(b_k^*, s^*(b_k^*))$ is the set of Nash equilibria of the two-stage game.

Figure 7c displays the equilibrium utility of Blue in the second-stage game for each of the three settings of Blue sensing (10, 30, and 50). This plot echoes what we had observed in the simultaneous choice game: High sensing is not played in equilibrium for any value of w between 0 and 10, and low sensing is preferred for most values of w. Figures 7a and 7b provide more insight by showing LR and probability of success of Blue in equilibrium. Here we see that, again, the choice of 50 for the sensing parameter is never preferred for either of these measures of effectiveness, and the choice of 10 results in considerably higher probability of success of Blue in equilibrium.

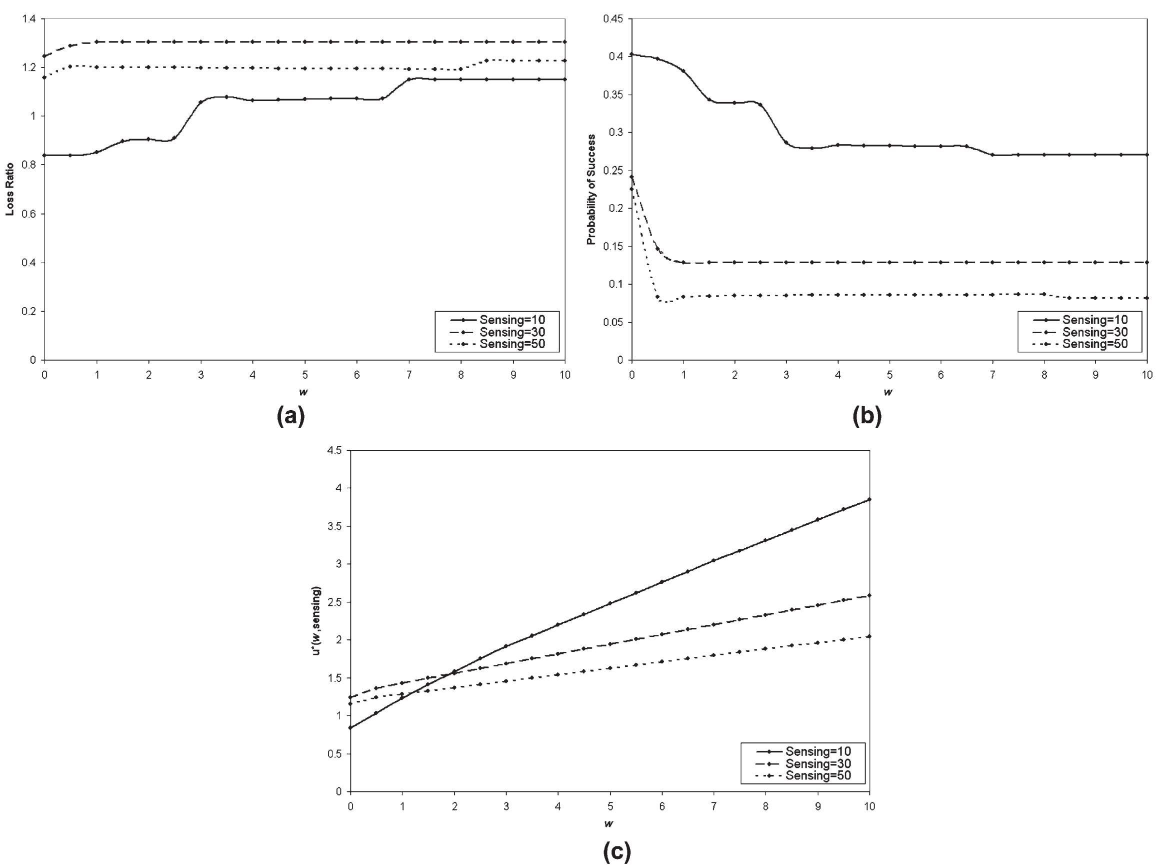

Figure 7. Equilibrium Loss Ratio (a), Probability of success (b), and utility (c).

Discussion of Results. The most surprising result of our analysis of this scenario is that improved sensing capability appears to hurt the Blue forces. Conventional wisdom would suggest that improved capabilities in general should improve outcomes, or at the very least not hurt the decision-maker. Although this intuition generally breaks down when rational self-interested agents engage in game-theoretic reasoning (Rubinstein, 1997), we do expect it to hold in zero-sum games since they can be modeled as optimization programs. To our bafflement, this intuition failed to hold in the ''Mosque'' scenario. (We get qualitatively similar results even when communication capacity is not a constraining factor in inter-squad communications.)

The first caveat that enters the picture indirectly is the implicit cost of medium and high sensing due to low target hit rate for distances between 10 and 20 cells. When sensing is low, the only agents detected are within a relatively high (60%) probability of hit range. When sensing is medium or high, some number of enemy troops appearing on an agent's situational awareness map are too far and unlikely to be hit (20% probability). Concentrating fire on agents that are far more likely to be hit will be beneficial when many agents are detected. On the other hand, added sensing capability will help Blue when only distant Red agents are detected. Thus, there is a clear tradeoff between the three sensing levels.

Our results also suggest that there is a cost to switching from a medium to a high sensing setting. To understand why it arises, we must delve precisely into the reason that the intuition that better capabilities should not be detrimental breaks down. The support of this intuition is the fact that the choices available to agents are still available when new capabilities are added, and this assumption breaks down in our case. When sensing capabilities are high, at any point during an identical run of the battle scenario, the situational awareness map of Blue may contain agents that would not have been detected when sensing was medium. Now suppose that Blue agents move toward or away from detected agents. When this is the case, agents that are not detected when sensing is medium would not influence the movement algorithm of Blue. If these agents are detected due to better sensing capability, Blue *no longer has the option of not reacting to agents that would not be on the map with a lower sensing setting*. The consequent tradeoff is due to the impulsiveness of Blue agents caused by the many sensing stimuli. In an attempt to react to sensing input, the agent may accidentally be driven to an undesirable state, which may have been avoided if the agent had an option not to react to a subset of the stimuli.

Thus, we have an apparent difficulty that we could not resolve because of the limitations of the MANA simulator. On the one hand, taking advantage of information requires strategies that prescribe the mode of behavior given information. On the other hand, when these are rule-based, they will tend to result in different sets of choices under the same actual combat states. In theory, we need strategies that prescribe possibly different choices of action from the same choice set for every possible SA map. In practice, such strategies would have to be extremely complex in order to produce interesting agent behavior. Indeed, this complexity will likely be confounding in actual combat scenarios as well. Certainly we cannot envision a commander even suggesting one of only two movement directions (left or right) for every possible scenario encountered, especially since the number of possible scenarios is likely to be extremely large. Instead, strategies do in practice have to be heuristic, and this will generally come at the cost of optimality that we theoretically can achieve.

''Mosque'' Scenario: Indirect Fire Support

This scenario is the same as the ''Mosque'' scenario just described, except an indirect fire squad is added to Blue that relies on communication with other Blue squads for its situational awareness. The experimental setup is unchanged, with three parameters (aggressiveness, cohesion, and sensing) available to Blue and two (aggressiveness and cohesion) available to Red. We approximate the game by taking a sample average of the utility function of Blue over 600 samples for each strategy profile.

Simultaneous Choice Model. Figures 8a and 8b depict equilibrium Loss Ratio and probability of

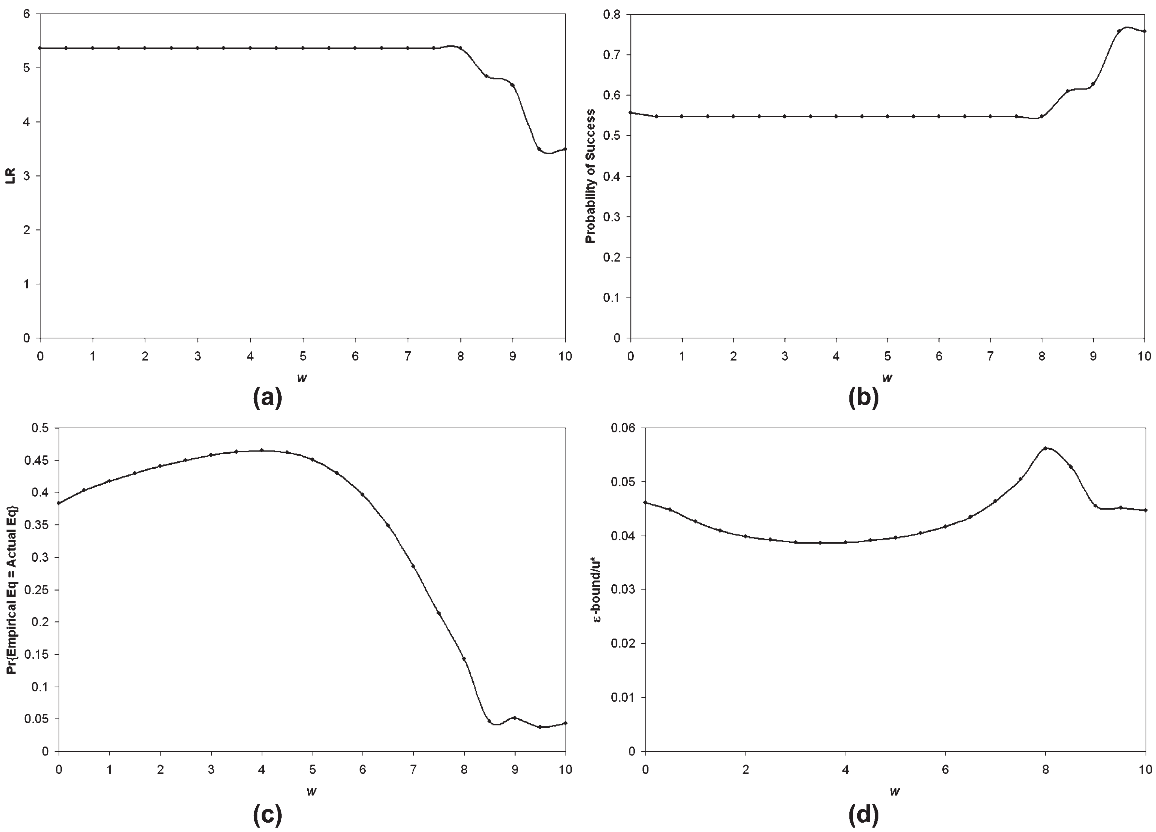

Figure 8. Loss Ratio (a), probability of success of Blue (b), probability that the equilibrium of the empirical game is an actual equilibrium (c), and maximum benefit to deviator with 95% confidence (d).

success of Blue in the simultaneous choice model. There are no real surprises, with probability of success increasing sharply when its importance is high ($w > 8.5$). The probability that our equilibrium analysis corresponds to an actual equilibrium of the game is shown in Figure 8c. Observe that for most values of w, our results have considerably more significance than for the Mosque scenario without indirect fire support. However, particularly when strategies appear to be in flux (w between 8.5 and 10), the probability drops significantly. Figure 8d shows the relative maximum benefit to deviation. We see that the values of the latter metric are low (between 3.5 and 6%) and vary little with w. Thus, it appears likely that our results provide a reasonable strategic guideline in this scenario.

We now look at the strategies that Blue and Red are expected to play in equilibrium. Sensing of Blue (Figure 9b) is 50 in equilibrium, unless probability of success is extremely important to Blue (high values of w), in which case the lowest sensing setting (sensing=10) is played. Compare this result with the previous scenario (no indirect fire support), where a sensing setting of 50 was never a part of the equilibrium strategy of Blue (in the range of w we present). However, as in the previous scenario, a sensing setting of 10 is the sole sensing setting in equilibrium when the probability of success has high importance. Figure 9a suggests that both Blue forces will be evasive for low values of w, but there will be some value of being aggressive occasionally when probability of success is important. Red forces will also be aggressive, maintaining this strategy for all values of w in the range (thus, we do not display it in a plot). Ignoring detected enemy troops entirely (aggressiveness = 0) is never a part of an equilibrium for either Blue or Red.

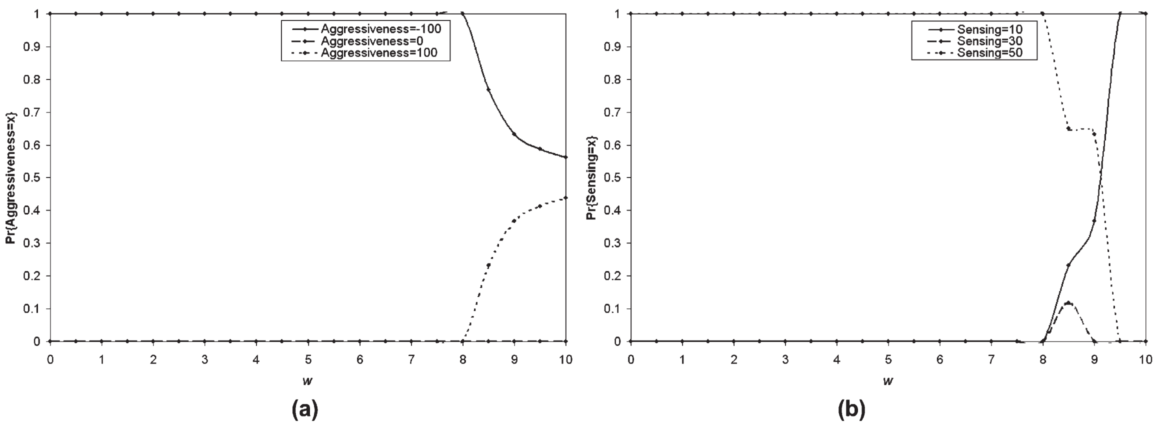

Figure 9. (a) Aggressiveness of Blue. (b) Sensing of Blue.

In Figure 10a, we observe that Blue will generally prefer to be spread out (cohesion=-100), unless the probability of reaching the goal (i.e., keeping Red leaders out of the mosque) is valued highly. In the latter case, Blue will form tight units (cohesion=100). Red (Figure 10b) will ignore the proximity of friendly forces for most values of w. Notable exceptions to this are w=0 (success of Red is irrelevant to Blue) and w between 9 and 10. In both cases, Red will intentionally attempt to spread out their troops (presumably to distract the Blue forces) with high probability.

Putting everything together, when w is low, Blue forces use high sensing, attempt to spread out, and avoid the enemy; Red forces are also spread out and avoid Blue. Thus, our results suggest that high sensing does indeed help Blue in avoiding the enemy troops. Since Red forces are spread out, Blue maximizes its LR by spreading out as well, as it has an advantage in one-on-one encounters. The fact that this response of Blue is optimal suggests that Blue would lose if it attempted to meet the spread-out Red forces with large groups, perhaps because this strategy would allow too many Red troops to escape unharmed into the mosque. Furthermore, in this scenario it appears that the advantage of the additional information about the location of enemy and friendly troops outweighs the implicit costs of sensing, such as lower hit rate and choice set dependence on information.

When probability of success is very important to Blue, the advantage of high sensing to Blue in improving LR is dissipated by the

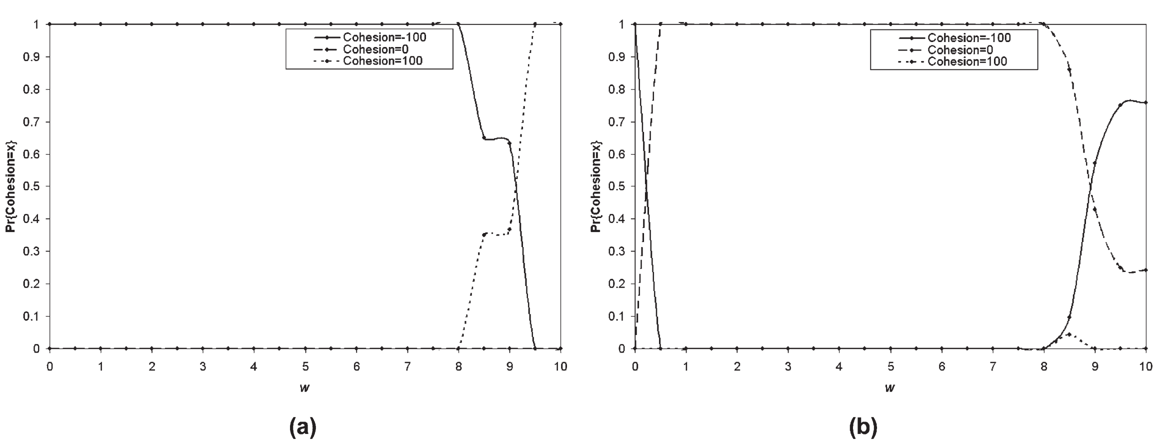

Figure 10. Cohesion of Blue (a) and Red (b).

apparent disadvantage in terms of probability of success. Blue will now mix between attacking and escaping the enemy upon encounter and will attempt to cluster its troops. Interestingly, the strategy of Red remains essentially the same. Thus, while clustering of Blue forces is detrimental to Blue's LR, it improves its probability of success. This is intuitive, since the success of Blue depends only on intercepting Red leaders on their way into the mosque. As there are relatively few of them, a close formation of Blue forces can ensure that they kill a detected Red leader without much disadvantage in detecting others. Given this strategy, however, higher sensing may be detrimental precisely because it will distract Blue clusters from their primary goal.

Sequential Choice Model. When we allow the Red player to observe Blue's choice of sensing parameter, the best choice of sensing for Blue is 50 for small to medium values of w, and 10 for large values of w (Figure 11c). In comparison, sensing=50 was never a good setting in the scenario in which Blue had no indirect fire support. On the other hand, sensing=10 still dominated for high values of w.

Figures 11a and 11b present LR and probability of success of Blue forces in equilibrium for each setting of the sensing parameter (10, 30, and 50). We observe that in this scenario, LR (in equilibrium) appears to be an increasing function of sensing, whereas probability of success is a decreasing function. The former result is different from what we observed in the first Mosque scenario, whereas the latter is identical.

That improved sensing results in higher LR comes as little surprise. When more Red troops are detected, the SA map is communicated to the indirect fire squad, and naturally this additional fire support results in considerably more Red casualties than Blue. What is surprising is

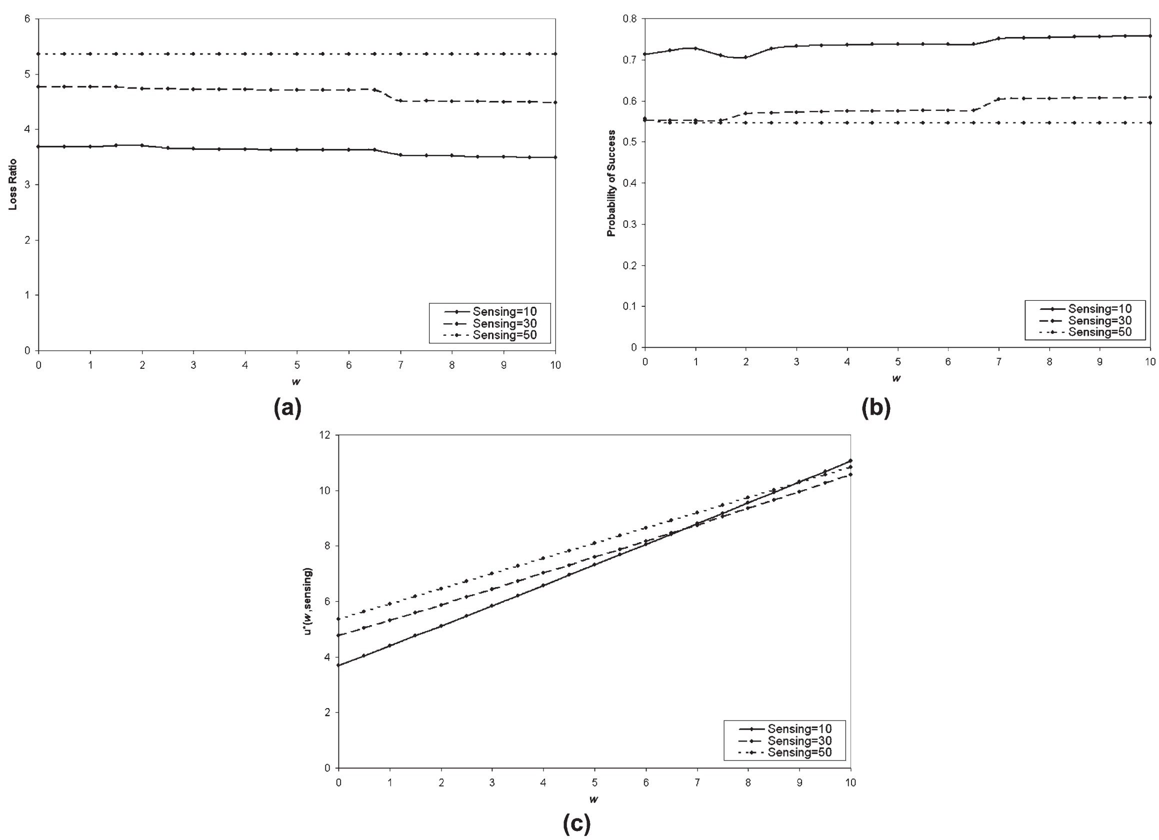

Figure 11. Loss Ratio (a), probability of success (b), and utility (c) of Blue in equilibrium.

that the same is not true of probability of success. From Figure 11b we see that low sensing results in a considerably higher probability of success than the other two settings. We conjecture that this is partly due to the distraction that the greater number of detected Red security forces produce: The indirect fire now has to share time between more Red forces, and there is therefore a greater chance that Red leaders will reach the mosque.

Discussion of Results. Table 2 summarizes our findings in the two Mosque scenarios. The most important differences in equilibrium strategies between the two scenarios are sensing of Blue and aggressiveness and cohesion of Red when w is low, and aggressiveness of Blue and cohesion of Red when w is high. The difference in sensing strategy is explained by its direct impact on the effectiveness of the indirect fire squad: greater sensing capabilities improve LR and make them preferred when the goal is relatively unimportant.

Explaining the reduced aggressiveness and cohesion of Red when Blue has indirect fire support is also relatively intuitive. With the indirect fire weapon present, whenever Red agents come in contact with Blue, they will face direct and indirect fire, the latter having potential effect on other Red agents that may be nearby. Similarly, Red must ensure that they are spread out enough that even if some are detected by the Blue, indirect fire weapon discharges hit only a relatively small proportion of Red agents, and will potentially have to share its firing time over spread out pockets of Red troops instead of concentrating fire on nearby locations. It is easy to see that concentrated fire would be more likely to be successful, since only one of the shots needs to hit the target. When probability of success is of greatest importance, Red will remain relatively isolated when indirect weapons are present, since spreading the indirect fire over a large region will increase the chance that Red leaders will sneak by unharmed into the mosque. Blue, on the other hand, will take advantage of the indirect fire and react to Red forces by sharing its time between engaging them (with the help of indirect fire) and escaping, thereby leaving the detected Red troops to the indirect fire squad and seeking out the other Red troops fleeing the scene.

CONCLUSION

Modern combat simulations can become a tremendous aid to military decision-making and have already been used in this capacity extensively. Yet, little combat simulation analysis work has been devoted to understanding optimal tactics of all adversaries, and not just the player that symbolizes the friendly force. When perfect intelligence regarding enemy tactics is available, game-theoretic analysis has little value, and simulations should be used instead to find the best response to known enemy plans. However, intelligence is typically at best uncertain and at worst unreliable, and a game-theoretic approach can significantly enhance the information available to a decision-maker, by either supporting the intelligence when the expected tactics of opponents are indeed rational, or casting doubt upon its

Table 2. Comparison of the two ''Mosque'' scenarios

	No indirect fire support		Indirect fire support	
	w=0	w=10	w=0	w=10
SIMULTANEOUS CHOICE				
Sensing (Blue)	medium	low	high	low
Aggressiveness (Blue)	evasive	usually passive	evasive	evasive or aggressive
Aggressiveness (Red)	passive	evasive	evasive	evasive
Cohesion (Blue)	isolated	clustered	isolated	clustered
Cohesion (Red)	indifferent	indifferent or clustered	isolated	usually isolated
SEQUENTIAL CHOICE				
Loss Ratio	high sensing	medium sensing	high sensing	medium sensing
Probability of success	low sensing	low sensing	low sensing	low sensing

reliability when the expected tactics of opponents are not rational.

The main contribution of this work was to develop a computational game-theoretic framework for military decision analysis based on combat simulations. Since we defined the interactions as a zero-sum game, we were able to use a general-purpose linear programming tool to solve games based on simulations. Extension of our analysis to general games is in principle straightforward, requiring merely the replacement of a zero-sum game solving tool with a general numerical equilibrium solver.

Random events are endemic to both battlefield scenarios as well as their simulations. To account for uncertainty in the simulation analysis due to sampling noise, we include a probabilistic evaluation of the results as a guidance regarding the likelihood that our results will hold if we had run the simulations an infinite number of times. Thus, even if simulation-based analysis fails to provide a conclusive answer, probabilistic analysis can provide additional guidance in the command and control decision-making process.

Discussion: Multiresolution Modeling

Often, the terms ''high-resolution'' and ''low-resolution'' in the context of ground combat simulations refer to the level of detail incorporated into the models of terrain effects on mobility and weapons systems. In addition, the increased role of networking and communication in combat has stimulated the need for modeling terrain effects on sensor and radio transmissions as a part of ''high-resolution'' investigations. Since increased resolution usually implies longer simulation time, a smaller space of tactical options is investigated, resulting in a tradeoff between model fidelity and the breadth of exploration of available options.

Phillips and Jackson (2005), as well as others (see NAS, 1997; Davis et al., 2001; and Bigelow and Davis, 2003) attempt to address this problem via multiresolution modeling (MRM). Proponents of MRM argue that faster, lower-fidelity tools should be run first to explore the space of options more thoroughly before high-resolution experiments are set up and run.

The work highlighted in this report lends itself to a series of three steps that can be followed as part of a multiresolution modeling and simulation (M&S) analysis effort. They are as follows:

1. Perform exploratory analysis to identify interesting factors that make up strategies of combatants. This could be done, for example, using agent-based tools.
2. Perform game-theoretic analysis that considers the restricted space of strategies identified in step 1 and determines equilibrium strategies.
3. Perform high-resolution simulation (of weapons effects, mobility, and communication) that uses the equilibrium strategies and replicates the scenario investigated with more fidelity.

NOTES

A. By strategy, we mean a decision by a player of any kind, whether that decision is tactical, operational, or strategic in the military sense. All that is important to us is that strategy choices of all players affect the ultimate outcomes.

B. In general, these will be mixed strategies, i.e., probability distributions over actions that players can take. Such actions in our setting are defined by the particular instantiations of strategy parameters.

C. For example, if victory is critical to both sides, the interaction is clearly zero-sum. Typically, however, there are many factors besides a one-time victory that are important, such as loss ratio, and sides may not agree on the relative importance of these, thereby eliminating the purely adversarial nature of the interaction. We assume here that Red and Blue agree on the relative importance of victory and loss ratio.

APPENDIX: THE LIMITATIONS OF MANA SIMULATION TOOL

It is extremely important to understand both the implications and the limitations of our analysis. An important limitation of any analysis of the nature we will present is the fact that any simulator is inherently imperfect. Thus, in order to draw conclusions from the analysis, we must understand the shortcomings of our simulator of choice. An important limitation of MANA is that it constrains the agent behavior to simple rules that are defined by parameter settings. It may seem especially strange to combine such rule-driven agents with a higher-level analysis that relies on the assumption of rationality of the players. We justify such an approach by drawing the distinction between players (Blue and Red) and individual agents in the simulations. We envision that the actual rational players are commanders, whereas the individual agents are soldiers. The orders from commanders to the soldiers are frequently rule-based (e.g., attack the enemy upon encounter if you are better equipped and not heavily outnumbered; withdraw otherwise), and there is relatively little room for soldiers to engage in game-theoretic analysis in the heat of the battle. Commanders, however, do have a certain bird's-eye view of the battle and will also need to plan the rules, training, and capabilities beforehand. Thus, the two-tiered approach that we take has an analogy in actual combat operations.

Nevertheless, we must still be cautious in using the results of our analysis to guide actual combat decisions. The first consideration is the resolution of the environment captured in MANA. MANA does allow agents to sense walls, foliage, and other simple obstacles that may impede movement and line of sight. However, not all weapons systems employed in MANA take into account the battlefield environment (e.g., weather or terrain), and, consequently, any sensitivity of weapon systems to the environment will not be captured within MANA.

ACKNOWLEDGEMENTS

The authors wish to acknowledge LTC Hattie D. Bouyer, Chief, Strategic Studies Division, U.S. Army Office of the CIO, whose sponsorship enabled these research efforts, and Stephanie Lonsinger, who gamely assisted in preparing the manuscript for publication.

REFERENCES

Bigelow, J., and P.K. Davis. 2003. *Implications for Model Validation of Multiresolution, Multiperspective Modeling (MRMPM) and Exploratory Analysis.* RAND Corporation, MR-1750-AF.

Davis, P., J.H. Bigelow, and J. McEver. 2001. *Exploratory Analysis and a Case History of Multiresolution, Multiperspective Modeling.* RAND Corporation, RP-925.

Davis, P. 2005. Personal Correspondence.

Galligan, D.P., M.A. Anderson, and M.K. Lauren. 2004. *MANA: Map Aware Non-Uniform Automota*, Users Manual. Defence Operational Support Technology Establishment.

Haywood, O.G. 1951. *Military Doctrine of Decision and the von Neumann Theory of Games.* RAND Corporation, RM-528.

Hamilton, T. and R. Mesic. 2004. *A Simple Game-Theoretic Approach to Suppression of Enemy Defenses and Other Time Critical Target Analyses.* RAND Corporation, DB-385-AF.

Hillestad, R.J. 1986. *SAGE: An Algorithm for the Allocation of Resources in a Conflict Model.* RAND Corporation, Working Draft.

Horne, G. 1997. *Data Farming: A Meta-Technique for Research on 21st Century Questions.* Briefing presented at the Naval War College, Newport, Rhode Island.

IBM Corporation. 2006. ILOG CPLEX webpage. http://www.ilog.com/products/cplex/

Ipekci, A.I. 2002. *How Agent Based Models Can Be Utilized to Explore and Exploit Non-Linearity and Intangibles Inherent in Guerilla Warfare.* Master's Thesis. Monterey, Calif.: Naval Postgraduate School.

Joint Concept Development and Revision Plan. 2004. http://www.dtic.mil/jointvision.

Kott, A, and M. Ownby. 2005. Tools for Real-Time Anticipation of Enemy Actions in Tactical Ground Operations. Presented at the 10th International Command and Control Research and Technology Symposium, McLean, Virginia.

Lachman, B., and R.J. Hillestad. 1986. *TAC-SAGE Reference Manual.* RAND Internal Note.

Lauren, M.K. 2001. Exploring the Value of Sensors to a RECCE Unit Using Agent-Based Models. *Journal of Battlefield Technology*, Vol 4, No 1.

Nash, J. 1951. Non-Cooperative Games. *Annals of Mathematics*, 286–295.

National Academy of Science (NAS). 1997. Technology for the United States Navy and Marine Corps: Becoming a 21st-Century Force. *National Academy Press*, Appendix E, Vol. 9, 2000-2035. Available at http://www.nap.edu/html/tech_21st/mse.htm.

Phillips, D., and J. Jackson. 2005. Using a Low Resolution Entity Level Modeling Approach. *PHALANX*, Vol 38, No 2, 15–17.

Porche, I., and B. Wilson. 2006. The Impact of Network Performance on Warfighter Effectiveness. RAND TR-329-A.

Porche, I., B. Wilson, and S. Witty. 2005. The Impact of Network Centric Operations on Warfighter Effectiveness: A Study Using Agent-Based Models. Presented at the 10th International Command and Control Research and Technology Symposium, McLean, Virginia.

Porche, Isaac, Bradley Wilson, and Susan Witty. 2007. The Impact of Network Centric Operations on Warfighter Effectiveness: A Study Using Agent-Based Models. *Military Operations Research*, Vol. 12, No. 3, 17–35.

Reeves, D., M.P. Wellman, J. MacKie-Mason, and A. Osepayshvili. 2005. Exploring Bidding Strategies for Market-Based Scheduling. *Decision Support Systems.* Vol 39, No 1, 65–85.

Rubinstein, A. 1997. *Models of Bounded Rationality.* MIT Press.

Schvartzman, J., and M.P Wellman. 2009. Stronger CDA Strategies Through Empirical Game-Theoretic Analysis and Reinforcement Learning. *Eighth International Joint Conference on Autonomous Agents and Multiagent Systems.* 249–256.

von Neumann, J. and O. Morgenstern. 1953. *Theory of Games and Economic Behavior.* Third Edition. John Wiley & Sons.

Vorobeychik, Y. 2009. Probabilistic Analysis of Simulation-Based Games. Technical Report. University of Pennsylvania. Available at http://www.seas.upenn.edu/~yev/2008/egbounds.pdf.

Walsh, W.E., R. Das, G. Tesauro, and J.O. Kephart. 2002. Analyzing Complex Strategic Interactions in Multi-Agent Systems. *AAAI-02 Workshop on Game Theoretic and Decision Theoretic Agents.* 109–118.

Wellman, M.P., A. Osepayshvili, J. MacKie-Mason, and D. Reeves. 2008. Bidding Strategies for Simultaneous Ascending Auctions. *The B.E. Journal of Theoretical Economics.* Vol 8, No 1.

Yiu, S.Y., A. Gill, and P. Shi. 2002. Investigating Strategies for Managing Civil Violence Using the MANA Agent Based Distillation. Presented at the Land Warfare Conference, Sydney, Australia.