Computational Models of Human Behavior in Wartime Negotiations

David V. Pynadath (pynadath@usc.edu), USC Institute for Creative Technologies

Ning Wang (ningwang@curiouslab.com), Curious Lab LLC

Stacy C. Marsella (marsella@ict.usc.edu), USC Institute for Creative Technologies

Abstract

Political scientists are increasingly turning to game-theoretic models to understand and predict the behavior of national leaders in wartime scenarios, where two sides have the options of seeking resolution at either the bargaining table or on the battlefield. While the theoretical analyses of these models is suggestive of their ability to capture these scenarios, it is not clear to what degree human behavior conforms to such equilibrium-based expectations. We present the results of a study that placed people within two of these game models, playing against an intelligent agent. We consider several testable hypotheses drawn from the theoretical analyses and evaluate the degree to which the observed human decision-making conforms to those hypotheses.

Keywords: Decision making; game theory; intelligent agents

Introduction

Political scientists are increasingly turning to computational models to understand and predict nations' wartime behavior (Fearon, 1995). Many such models combine both military and political processes, where battlefield decisions occur within the context of an overall negotiation over a contentious resource (Filson & Werner, 2002; Smith & Stam, 2004). Game-theoretic models of these processes seek to capture possible outcomes on the battlefield and at the bargaining table (Powell, 2001, 2004; Slantchev, 2003). These models hypothesize equilibrium strategies that correspond to the behaviors of nations in real-world scenarios.

While the theoretical analyses of such game-theoretic models is suggestive of their representational power, it remains an open question as to how well they capture actual human behavior in wartime negotiation. These models focus on equilibrium behavior, where both sides optimize their outcomes in response to the others' behaviors. The computational challenges of such optimization often require the equilibrium analyses to make simplifying assumptions (e.g., to reduce uncertainty about the opponent). However, people are not constrained to adopt these same assumptions when making their decisions, so it is possible that human behavior in the face of this uncertainty may greatly deviate from the predictions generated by such theoretical models.

On the other hand, these computational models easily lend themselves to an experimental setting, where we can pit a human player against an intelligent agent playing according to the model of interest. In other words, we can place human players within the game hypothesized by a model and have them negotiate with an agent. We can then observe human behavior and quantify the degree to which that behavior conforms to the expectations generated by the model.

This paper presents a human subject study where we implemented games corresponding to two wartime negotiation models from the literature (Powell, 2004; Slantchev, 2003). We present behavior hypotheses extracted from the theoretical analyses of these models. We analyze the observed human behavior to see that it generally satisfies these hypotheses. However, there are also interesting deviations from these theoretical expectations that suggest possible extensions to the models to better capture human decision-making.

Wartime Negotiation Models

A number of formal models in the literature represent war as a costly process embedded within a negotiation game. In these models, two sides are in a dispute over a desirable resource, such as territory claimed by both sides. The game begins with some initial split of the territory. The game progresses round by round, with each round consisting of one side proposing a split of the territory, the other side responding to that proposal, and a possible battle. The game ends with a final split achieved by either an agreement on the proposed split or a decisive military victory by one side on the battlefield.

To facilitate a game-theoretic analysis, these models make simplifying assumptions regarding military outcomes. In particular, the probabilities associated with the battlefield are fixed, so that one side's probability of winning does not change during the course of the game, regardless of previous military outcomes. The costs of a single battle are also fixed throughout the course of the game. In our study, we present these costs to the human players in terms of troops lost.

A critical property of these models is uncertainty about the likelihood of battlefield outcomes. If both sides had complete information about their probability of winning battles, they could do an exact cost-benefit analysis and immediately agree upon an acceptable territorial split. The models we consider instead have incomplete information, where one side is ignorant of the probability of battlefield outcomes. As the war progresses, this side will gain information by observing battle and bargaining outcomes, re-evaluate its prospects, and make different decisions on offers and battles. This asymmetry lends itself to our human subject study, as we can give the agent complete information about the game probabilities, but hide that information from the human player. Our experiments will then allow us to study how people update their beliefs based on the information that is revealed in the game.

We chose two models (Powell, 2004; Slantchev, 2003) for this investigation, based on their impact on the field and their appropriateness for a human-agent game interaction.

Powell Model

The Powell model proceeds as follows (Powell, 2001, 2004):

1. Player 1 makes an offer of $x\%$ of the territory.

2. Player 2 decides to accept, reject, or go to war.

   (a) If acceptance, Player 2 gets $x\%$ of the territory, Player 1 gets $(100-x)\%$, and the game is over.
   (b) If war, both sides incur the battle costs. Player 1 collapses with probability p_1 and Player 2 collapses with probability p_2.
      i. If only Player 1 collapses, Player 2 gets all of the territory, and the game is over.
      ii. If only Player 2 collapses, Player 1 gets all of the territory, and the game is over.
      iii. Otherwise, return to Step 1.
   (c) If rejection, Player 1 decides whether or not to go to war.
      i. If war, a battle occurs exactly as in Step 2b.
      ii. If not, return to Step 1.

The following properties distinguish this model from the Slantchev model (described in the next subsection):

Battle: Attacking is a choice (if an offer is rejected).

War state: In each battle, there is a fixed probability that you win (lose) the overall war and gain (lose) all of the territory.

Counteroffers: There are no offers made by Player 2.

Player 1, the offering side, does not know the probabilities of collapse (p_1 and p_2), but Player 2 *does* know these probabilities. Thus, Player 1 is uncertain about the two sides' relative military strength and, consequently, the feasible agreements. The equilibrium behavior can be described as *screening*, where Player 1 will make a series of increasingly attractive offers, expecting weaker opponents to accept early in the process, thus screening them out before making the higher offers necessary to appease stronger opponents (Powell, 2004).

Slantchev Model

The Slantchev model includes an additional variable, *military position* (in $\{0, 1, 2, \ldots, N\}$), that represents the relative gains made by the two sides in the war so far (Slantchev, 2003). The game under this model proceeds as follows:

1. The initiating player makes an offer of $x\%$ of the territory.

2. The responding player decides to accept or reject the offer.

   (a) If acceptance, the responding player gets $x\%$ of the territory, the initiating player gets $(100-x)\%$, and the game is over.
   (b) If rejection, continue to Step 3.

3. Battle occurs, and both sides incur the fixed costs. Player 1 wins the battle with probability p, Player 2 with probability $1-p$.

   (a) If Player 1 wins, military position increases by 1. If it reaches N, then Player 1 receives all the territory and the game is over.
   (b) If Player 2 wins, military position decreases by 1. If it reaches 0, then Player 2 receives all the territory and the game is over.

4. Return to Step 1 with initiating and responding players reversed.

This model deviates from Powell's as follows:

Battle: There is a battle every round.

War state: A single battle does not directly end the war, but affects the military position variable. Collapse occurs only if military position hits its maximum or minimum value.

Counteroffers: Both sides alternate in making offers.

Like the Powell model, Player 1 does not know the battle probability (p), but Player 2 does, so the equilibrium behavior again exhibits some screening. However, the Slantchev model provides Player 1 with the additional information source of Player 2's counteroffers. Furthermore, the military position provides another variable for the sides to consider, in that their offering behavior will change depending on which side is in a stronger position in the overall war.

PsychSim Agents

We implemented both the Powell and Slantchev games within PsychSim, a multiagent framework for social simulation (Marsella, Pynadath, & Read, 2004; Pynadath & Marsella, 2005). PsychSim agents have their own goals, private beliefs, and mental models about other agents. They generate their beliefs and behaviors by solving partially observable Markov decision problems (POMDPs) (Kaelbling, Littman, & Cassandra, 1998), whose quantitative transition probabilities and reward functions are a natural fit for the game-theoretic dynamics of our chosen models of wartime negotiation.

PsychSim agents have a *theory of mind* that allows them to recursively model other agents (e.g., their beliefs, rewards, etc.), form expectations about their behavior, and choose optimal actions as a best response. With sufficient computation, the PsychSim agent's optimal action corresponds to the equilibrium strategy. However, we can also limit the agent's horizon when computing expected values and the depth of recursion in modeling others. By doing so, the agent can quickly compute approximate best-response actions even when a human opponent deviates from the equilibrium.

The behavior of the PsychSim agents in both the Powell and Slantchev models roughly corresponds to the informed side's equilibrium strategy. For example, when starting with less territory in the Powell model, the PsychSim agent chooses war until its opponent makes an offer that exceeds its threshold of acceptability, computed as a function of the probability of military collapse. Under the Slantchev model, the agent also rejects any offer below a threshold, but that threshold changes based on the current military position. In particular, if the military position favors the agent's side, the threshold is higher than it would be otherwise. The counteroffers made by the agent are similarly lower when the military position is in its favor than they would be otherwise.

While both Powell and Slantchev focused on the case where the uninformed side also started as the satisfied side, we can also model the case where the uninformed side starts as the *dissatisfied* side. We change the initial distribution of territory from having the human player start with 72% of the territory as the satisfied side, to having the human player start with only 28%. The PsychSim agent computes its policy of

behavior using the same algorithm in both cases, although the resulting strategy is slightly different. Under Powell, the agent playing the satisfied side (with 72% of the territory) will no longer attack when receiving an unacceptable offer. Instead, it will simply reject the offer, hoping to avoid a battle that will risk collapse and loss of all its territory. Under Slantchev, the agent's thresholds as the satisfied side will be universally higher than those as the dissatisfied side.

Wartime Negotiation Study

We used these agents in a study of how people make decisions in wartime negotiation games. The subjects played each model twice: once as the satisfied side (starting with 72% of the territory) and once as the dissatisfied side (starting with 28%), leading to four experimental conditions: Powell72, Powell28, Slantchev72, and Slantchev28. For each condition, the subject played as Player 1 against a PsychSim agent until the two sides agreed on a split or one side achieved a military victory. If neither occurred within 15 rounds, the game terminated with the sides staying at the initial division of territory. In the two Powell conditions, both sides have the same probability of collapse ($p_1 = p_2 = 0.1$). In the two Slantchev conditions, we use the same probability of winning for Player 1 ($p = 0.3$) and the same initial military position, with Player 1 slightly closer to losing the war (3 on a range from 0 to 10).

Hypotheses

The Powell and Slantchev models yield hypotheses about behavior that we might see in our human subject data:

Screening Behavior   The uninformed side tries to find the minimal offer that is acceptable to its opponent. It does so by progressively increasing its offer until its opponent accepts, gradually screening out weaker opponents who are willing to accept lower offers. We expect to see players make these increasing offers under both Powell and Slantchev models.

Principle of Convergence   Warfare ceases to be useful when there is no information to gain, at which point the sides can both agree on a settlement. Given the static battle probability and the lack of signaling moves in the Powell model, the potential information gain should be exhausted sooner than in the Slantchev model. As a result, we would expect settlement to be reached sooner under the Powell model, where the only information gain is through rejected offers.

Information Asymmetry   Because of the screening behavior, Slantchev claimed, "as war progresses, the outcome becomes less advantageous for the worse informed party." We thus expect settlements that take more rounds to be less favorable to the human players. Furthermore, we expect this trend to be more pronounced under Powell, where the players receive less information than they do under Slantchev.

Total victory   Total military victory (i.e., one side winning all of the territory on the battlefield) is rare, as war typically reveals information quickly enough for both sides to reach

settlements instead. The possibility of collapse in a single battle under the Powell model should make total victory much more common. We would thus expect that negotiated settlements (as opposed to total victories) to be *less* common under the Powell model than under Slantchev.

War avoidance   Both models provide incentives for sides to sacrifice territory to avoid a costly battle. We would expect players who expressed a more positive attitude toward war (ATW) to exhibit more of a willingness to engage in war and give up less territory in the final settlement.

Military Asymmetry   If the uninformed side is also at a military disadvantage (as is the case for our human players), then we expect it to overestimate its probability of winning the war, thus making lower offers than it would make in the complete-information case. Therefore, we would expect to see lower offers when the players receive less information about their relative military strength (i.e., under Powell) than when they receive more (i.e., under Slantchev).

Battle outcomes   A battle in the Slantchev model will make the victor more optimistic and more willing to delay agreement. We would thus expect players to make lower offers after winning a battle than they would after losing a battle.

Starting Territory   Our four experimental conditions could engender different reference points (Neale & Bazerman, 1991) when people play as the satisfied or dissatisfied sides (starting with 72% or 28% of the territory, respectively). We hypothesize that satisfied sides will make fewer concessions, as the *endowment effect* makes players less willing to give up territory already owned (Kahneman, Knetsch, & Thaler, 1991). Thus, we expect a player to offer less when starting with more territory than when starting with less territory. Similarly, we expect the dissatisfied side will end up with less territory, because any territory gain, however small, would be more likely considered as satisfactory. Furthermore, because the satisfied side has more to lose through a military outcome, we expect that the *difference* between initial and final territorial splits will be more favorable for the dissatisfied sides.

Study Population

We recruited 240 participants, of an average age of 35, via Amazon Mechanical Turk. 51% of the participants were female, and 49% were male. 65% of the participants were from the United States, 29% from India, and 6% from other countries. Regarding the participants' highest level of education, 12% of the participants had some high school or high school diploma, 63% had some college or college degree, and 25% had some graduate school or graduate degree. 13% of the participants used a computer for 1-4 hours a day, 43% for 5-8 hours a day, and 44% for more than 8 hours a day.

Procedure

After being assigned an anonymous ID, each participant read an information sheet about the study and then filled out a background survey. Next, the participant played the negotia-

tion game four times, each time against a different agent from one of the conditions. The order of the four agents the player negotiated with was randomized. During the negotiation, the participant filled out an in-game survey. Following each negotiation game, the participant filled out an opinion survey. We designed the study to be completed within an hour, and the average duration of the study was 32 minutes.

Measures

Background Survey The background survey asked about the participant's age, gender, nationality, education, computer experience, Attitude Towards War (Dupuis & Cohn, 2006), Social Orientation (Van Lange, De Bruin, Otten, & Joireman, 1997) and attitude towards Inappropriate Negotiation (SINS, from (Robinson, Lewicki, & Donahue, 2000)).

Opinion Survey The opinion survey contained questions regarding the participant's goals during the game, and questions from the Subjective Value Index (SVI) survey on opinions about oneself, the negotiation outcome, the process and the opponent (Curhan, Elfenbein, & Xu, 2006).

In-Game Survey The in-game survey asked participants to predict the opponent's responses. For example, after making an offer in the Powell game, participants said whether they expected their opponent to accept the offer, reject it or attack.

Game Logs The game logs captured the participant's actions, the PsychSim agent's actions and the world states (e.g. number of troops, military position, and territory).

Results

We had 240 participants in the study. Each participant played four different games, one under each condition. Data from incomplete games were discarded. In the end, we had 238 games in the Powell72 condition and 239 games in the Powell28, Slantchev72 and Slantchev28 conditions.

Hypothesis: Screening Behavior

This hypothesis states that the participants' behavior is most likely to be screening, by making incrementally higher offers to find out the lowest offer that satisfies the opponent. We analyzed the dynamics of the participants' offers to see whether they increased, decreased, or stayed the same from one round to the next. The results of Figure 1 show that, by and large, the human players exhibit screening behavior (i.e., more increases than the alternatives).

Hypothesis: Principle of Convergence

The hypothesis states that under the Powell model, settlement should be reached sooner, compared to the Slantchev model. We compared the number of rounds it took for both sides to reach an agreement under these two models, excluding games that ended with one side winning the war (instead of reaching an agreement). Results show that, contrary to the hypothesis, it took participants significantly more rounds to reach an agreement when interacting with the Powell model
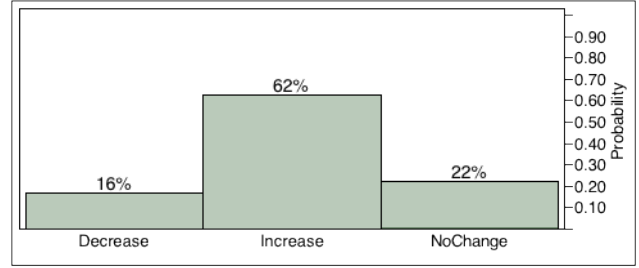

Figure 1: Distribution over participants' offer dynamics.

than when interacting with the Slantchev model ($p < .0001$, Mean$_{Powell}$=3.13, Mean$_{Slantchev}$=1.95).

Hypothesis: Information Asymmetry

This hypothesis predicted that settlements taking more rounds to reach agreement would be less favorable toward the human player, with the effect being more pronounced in the Powell model. Excluding games with no agreement, we analyzed the correlation between the territory participants ended up with and the total number of rounds needed to reach the settlement. In general, there is a weak yet significant negative correlation between the territory participants got in the settlement and the number of rounds needed to reach that settlement ($r = -.1965$, $p < .0001$), providing evidence in favor of this hypothesis. The correlation is of medium strength both in the Slantchev ($r = -.2372$, $p < .0001$) and Powell games ($r = -.2309$, $p = .0004$), failing to demonstrate the hypothesized difference between the two models.

Hypothesis: Total Victory

Under this hypothesis, we expect that negotiated settlements to be less common under Powell than under Slantchev because of the possibility of immediate collapse in the former. The data bore out this hypothesis, as fewer games in the Powell model ended in a settlement than in the Slantchev model ($p < .0001$, Mean$_{Powell}$=48.0%, Mean$_{Slantchev}$=60.8%). It is also interesting to observe that settlements were much rarer when the player started with 28% territory than when starting with 72% ($p < .0001$, Mean$_{28}$=42.77%, Mean$_{72}$=66.25%).

Hypothesis: War Avoidance

We hypothesized that players who are more pro-war would be less willing to give up territory, and more willing to go to war. We measured the participants' attitudes towards war (ATW) in the background survey, where higher ATW scores indicate more of a pro-war attitude, and lower ones indicate an anti-war attitude. We did not find significant correlations between ATW and the average offers participants made ($r = -.0172$, $p = .5950$). We also did not find a correlation between ATW and the number of rounds played in the game ($r = .0220$, $p = .4979$). Battles were not a choice in the Slantchev model. In the Powell72 condition, the human player never initiated an attack, because the agent would always do so first. Therefore, we analyzed the correlation between ATW and the number of player-initiated attacks in only the Powell28 condition and
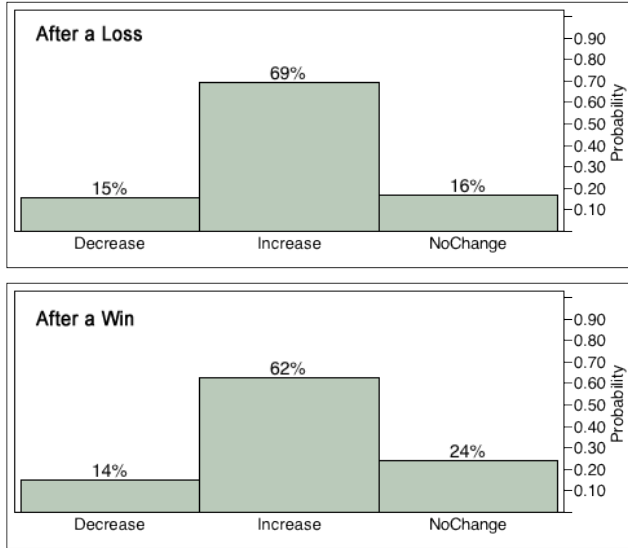
Figure 2: Change of offers made after battle outcomes.

found a marginally significant weak correlation ($r = .1107$, $p = .0877$). Thus, there was only the slightest of evidence in favor of his hypothesis.

Hypothesis: Military Asymmetry

This hypothesis states that in the Powell model, where less information is revealed to the players, the players will make lower offers than they would under Slantchev. However, when interacting with the Powell model, participants made slightly *higher* offers than when they interacted with the Slantchev model ($p < .0001$, Mean$_{Powell}$=36.26, Mean$_{Slantchev}$=33.29), exactly the opposite of our hypothesis.

Hypothesis: Battle Outcomes

We expect players to make lower offers after winning a battle than they would after losing. We ignore the Powell model, which gives players no information about battle outcomes (beyond game-ending collapses). Under Slantchev, players lost 68% of the battles and won only 32%. When players won a battle, the offers they made next were significantly lower than when they lost ($p = .0441$, Mean$_{Won}$=31.70, Mean$_{Lost}$=35.88), thus bearing out the hypothesis.

Hypothesis: Starting Territory

Initial Offer   We compared the offers that participants made at the beginning of the negotiation. Participants made significantly higher offers when starting with 28% territory compared to 72% ($p < .0001$, Mean$_{28}$=34.20, Mean$_{72}$=27.95).

Average Offer   We compared the average offers the participants made during the negotiation. When starting with 28% territory, participants made significantly higher offers than with 72% ($p < .0001$, Mean$_{28}$=38.84, Mean$_{72}$=30.71).

End Territory   We compared the percentage of territory the participant had when the game ended. When starting with 28% territory, participants ended up with signifi-
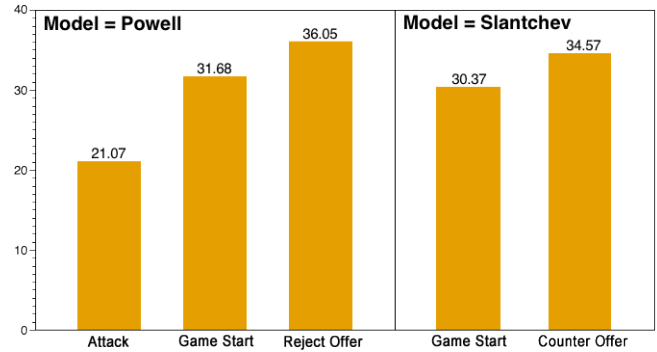

Figure 3: Participant offers in response to opponent actions.

cantly less territory than when starting with 72% ($p < .0001$, Mean$_{28}$=22.31, Mean$_{72}$=45.55).

Net Territory Gain   Beyond the impact on absolute territory, we also hypothesized that the starting territory would affect the *relative* gain/loss in territory from the beginning to the end of the game. When starting with 28% territory, participants lost significantly less territory than when they starting with 72% territory ($p < .0001$, Mean$_{28}$=-5.69, Mean$_{72}$=-26.45). Thus, the observed behavior conformed to all of our expectations about the effect of the initial division of territory.

Offers In Reaction to Opponent Actions

We also compared the participants' offers under both models in reaction to their opponent's actions. In the Powell model, when not accepting the participant's offer, the opponent chose to either simply reject the offer or to attack the participant. In the Slantchev model, attacking was not a choice, but the opponent did make counteroffers when not accepting the participant's offer. The differences were significant when interacting with the Powell model ($p < .0001$) and Slantchev model ($p < .0001$). The results of Figure 3 show that attacking the participants prompted lower offers, while less aggressive actions (e.g. rejecting without attacking, or making a counteroffer) typically resulted in higher offers.

Discussion

As we can see from the previous section, much of the observed behavior conformed to the expectations generated by the theoretical analyses of the Powell and Slantchev models. The information asymmetry that is critical to both models had the expected impact on the human players, as they clearly exhibited the hypothesized screening behavior. Furthermore, we also observed direct evidence of the claim that "as war progresses, the outcome becomes less advantageous for the worse informed party" (Slantchev, 2003). Our agent-based experimental setup also allowed us to try a variation of the game on the starting territory, and the data provided strong evidence for our hypotheses regarding that variation.

However, there were also some interesting deviations from our hypothesized behavior. Slantchev's *Principle of Convergence* hypothesized that warfare ceases to be useful when there is no information to gain. Our derived hypothesis

viewed the Powell game as more information-poor than the Slantchev one, leading us to expect that settlements would be reached sooner in the former. However, our data showed the opposite trend. We can at least partly attribute this deviation to people assessing the information gains differently than was prescribed in the two models' equilibrium analysis. Under the Powell model, players may try to gather more observations of their opponents' behavior, because they do not realize that the agent's threshold of acceptance does not change. Conversely, under the Slantchev model, the players may feel overly certain in their beliefs upon observing only a few (or even only one) counteroffers from their opponent.

We also saw little evidence to support our *War Avoidance* hypothesis. We did not find any link between a participant's attitude towards war and the offers they made, the duration of the war, nor the number of attacks. It is likely that some participants did not carry their attitudes toward war over into this abstract game setting. However, we also need to further differentiate within the ATW scale about why people are anti-war and what types of war they are against.

The deviation from the *military asymmetry* hypothesis is harder to explain. From the very beginning of a Slantchev game, the players can observe that their military position puts them closer to losing the war than winning the war. As the game progresses, they can potentially observe that the battle probability is not in their favor. We would thus expect players to be more pessimistic about their chances under the Slantchev model and, thus, to make higher offers to quickly appease their opponents. However, players made higher offers in the Powell model, where there was no feedback about the war outcomes, nor the other side's valuations.

Therefore, we need a more fine-grained analysis of when players in Slantchev made the unexpectedly low offers. For example, the data supported our *Battle Outcomes* hypothesis, where players made lower offers after winning a battle. By isolating such cases, we may see that the general *military asymmetry* hypothesis holds, but we can also understand the in-game contingencies that would override the general trend.

Furthermore, despite the general conformity over all of the data, not *every* player's behavior conformed to our hypotheses. We must therefore analyze the data to identify the more fine-grained contingencies and the individual differences among our participants. Such an analysis would give us a better understanding of how the participants viewed, for example, the potential information received in the game, their possible military outcomes, etc. This analysis will also guide future studies by suggesting further instrumentation that is needed to gather the required in-game data.

With these additional analyses and data, we can build upon the field's game-theoretic models to develop higher-fidelity models of human behavior in such wartime negotiation scenarios. Ideally, these models will help bridge the gap between the theoretical computational frameworks and the decision-making we see in the real world. By doing so, they will provide an invaluable computational tool for political scientists to explore a richer set of contingencies and individual differences and hopefully provide better predictions and explanations of behavior in wartime negotiation.

Acknowledgments

References

Curhan, J. R., Elfenbein, H. A., & Xu, H. (2006). What do people value when they negotiate? Mapping the domain of subjective value in negotiation. *J. of Personality and Social Psychology*, *91*, 493–512.

Dupuis, E. C., & Cohn, E. S. (2006). A new scale to measure war attitudes: Construction and predictors. *J. of Psychological Arts and Sciences*, *3*(1), 6–15.

Fearon, J. D. (1995). Rationalist explanations for war. *International Organization*, *49*, 379-414.

Filson, D., & Werner, S. (2002). A bargaining model of war and peace: Anticipating the onset, duration, and outcome of war. *American J. of Political Science*, *46*, 819-838.

Kaelbling, L. P., Littman, M. L., & Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, *101*, 99-134.

Kahneman, D., Knetsch, J. L., & Thaler, R. H. (1991). Anomalies: The endowment effect, loss aversion, and status quo bias. *The J. of Economic Perspectives*, *5*(1), 193-206.

Marsella, S. C., Pynadath, D. V., & Read, S. J. (2004). PsychSim: Agent-based modeling of social interactions and influence. In *ICCM* (p. 243-248).

Neale, M. A., & Bazerman, M. H. (1991). *Cognition and rationality in negotiation*. Free Press.

Powell, R. (2001). *Bargaining while fighting*. Doctoral dissertation, University of California, Berkeley.

Powell, R. (2004). Bargaining and learning while fighting. *American J. of Political Science*, *48*(2), 344–361.

Pynadath, D. V., & Marsella, S. C. (2005). PsychSim: Modeling theory of mind with decision-theoretic agents. In *IJCAI* (p. 1181-1186).

Robinson, R. J., Lewicki, R. J., & Donahue, E. M. (2000). Extending and testing a five factor model of ethical and unethical bargaining tactics: Introducing the SINS scale. *J. of Organizational Behavior*, *21*, 649–664.

Slantchev, B. L. (2003). The principle of convergence in wartime negotiations. *American Political Science Review*, *97*, 621–632.

Smith, A., & Stam, A. C. (2004). Bargaining and the nature of war. *The J. of Conflict Resolution*, *48*(6), 783-813.

Van Lange, P. A. M., De Bruin, E. M. N., Otten, W., & Joireman, J. A. (1997). Development of prosocial, individualistic, and competitive orientations: Theory and preliminary evidence. *J. of Personality and Social Psychology*, *74*(4), 733–746.