



January 2007

Socio-Cultural Games for Training and Analysis

Barry G. Silverman

University of Pennsylvania, basil@seas.upenn.edu

Gnana K. Bharathy

University of Pennsylvania

Michael Johns

University of Pennsylvania

Roy J. Eidelson

University of Pennsylvania

Tony E. Smith

University of Pennsylvania

See next page for additional authors

Follow this and additional works at: https://repository.upenn.edu/ese_papers

Recommended Citation

Barry G. Silverman, Gnana K. Bharathy, Michael Johns, Roy J. Eidelson, Tony E. Smith, and Benjamin Nye, "Socio-Cultural Games for Training and Analysis", . January 2007.

Postprint version. To be published in *IEEE Transactions on Systems, Man, and Cybernetics*, October 2007, 49 pages.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/ese_papers/301
For more information, please contact repository@pobox.upenn.edu.

Socio-Cultural Games for Training and Analysis

Abstract

This paper presents a theory for role playing simulation games intended to support analysts (and trainees) with generating and testing alternative competing hypotheses on how to influence world conflict situations. Simulated leaders and followers capable of playing these games are implemented in a cognitive modeling framework, called PMFserv, which covers value systems, personality and cultural factors, emotions, relationships, perception, stress/coping style and decision making. Of direct interest, as Section 1.1 explains, is codification and synthesis of best-of-breed social science models within PMFserv to improve the internal validity of the agent implementations. Sections 2 and 3 present this for leader profiling instruments and group membership decision-making, respectively. Section 4 then offers two real world case studies (The Third Crusade and SE Asia today) where the agent models are subjected to Turing and correspondence tests under each case study. In sum, substantial effort on game realism, best-of-breed social science models, and agent validation efforts is essential if analysis and training tools are to help explore cultural issues and alternative ways to influence outcomes. Such exercises, in turn, are likely to improve the state of the science as well.

Keywords

Leaders and followers, spread of ideas, strategy games, personality and culture, agent-based simulation

Comments

Postprint version. To be published in *IEEE Transactions on Systems, Man, and Cybernetics*, October 2007, 49 pages.

Author(s)

Barry G. Silverman, Gnana K. Bharathy, Michael Johns, Roy J. Eidelson, Tony E. Smith, and Benjamin Nye

Socio-Cultural Games for Training and Analysis

Barry G. Silverman¹, Gnana Bharathy¹, Michael Johns¹, Roy J. Eidelson², Tony E. Smith¹, Benjamin Nye¹

1- Electrical and Systems Engineering Dept.

2- Asch Center for EthnoPolitical Conflict

University of Pennsylvania, Philadelphia, PA 19104-6315

barryg@seas.upenn.edu

March 2007

This paper presents a theory for role playing simulation games intended to support analysts (and trainees) with generating and testing alternative competing hypotheses on how to influence world conflict situations. Simulated leaders and followers capable of playing these games are implemented in a cognitive modeling framework, called PMFserv, which covers value systems, personality and cultural factors, emotions, relationships, perception, stress/coping style and decision making. Of direct interest, as Section 1.1 explains, is codification and synthesis of best-of-breed social science models within PMFserv to improve the internal validity of the agent implementations. Sections 2 and 3 present this for leader profiling instruments and group membership decision-making, respectively. Section 4 then offers two real world case studies (The Third Crusade and SE Asia today) where the agent models are subjected to Turing and correspondence tests under each case study. In sum, substantial effort on game realism, best-of-breed social science models, and agent validation efforts is essential if analysis and training tools are to help explore cultural issues and alternative ways to influence outcomes. Such exercises, in turn, are likely to improve the state of the science as well.

Keywords: Leaders and followers; spread of ideas; strategy games; personality and culture, agent-based simulation

I. INTRODUCTION AND PURPOSE

Gaming and simulation of socio-cultural groups is a newly evolving field, motivated by the need to better understand how leaders and followers behave, what motivates them, how dangerous ideas spread, and how they might be influenced to cooperate, mitigate conflicts, and benefit the overall good. Green and Armstrong (2003) study the array of methods for forecasting conflict and show that predictions are significantly improved when

subjects first participate in role playing games about the issues at stake. Hence, one aim of this research is to isolate the components needed for a generic role playing game to be used to rapidly mock up a class of conflicts commonly encountered in today's world. In other words, create a widely applicable game generator. Since it is often impossible to find humans to play all the roles of such games, or to play out all the possible scenarios, a second aim is to create plausible models of leaders and followers based on first principles about what makes them "tick" and so they may play some of the roles in the game. If these cognitive agents are realistic, they can help trainees and analysts explore the range of their possible actions under varieties of conditions, thereby helping others to see more clearly how to influence them and elicit their cooperation. A related assumption based on evidence from video- and multi-player online-games, is that if the agents have sufficient realism, that should further motivate players (trainees) to be engaged and immersed in role playing games or online interactive scenarios. The benefits of these first two goals lie in encouraging trainees and analysts to arrive at explanations of what is going on in a situation.

A 'catch-22' of these two goals is that, agent based simulation games will be more valuable the more they can be imbued with realistic leader and follower behaviors, while the social sciences that can reliably contribute to this undertaking are made up of many fragmented and narrow specialties, and few of their 'models' have computational implementations. The third aim of this research is thus to improve the science by synthesizing relevant first principles and best of breed social science models so they can be tested in agent-based games, and by that to cause the limitations of these models to be exposed and improved.

It is a human tendency to project our own value systems upon others and presume they want the same things we want (the mirror bias). Once we form such hypotheses, we tend to look only for confirming evidence and ignore disconfirming facts (the confirmation bias). Heuer (1999) points out that it is vital to break through these and related biases, and that methodical approaches such as realistic simulations, if well done, might help to elucidate and explore alternative competing hypotheses of agents' motivations, intentions and consequent behavior. Thus, the fourth aim is to generate and test new hypotheses, which is also a potential benefit of simulations.

We turn to a review of some of the literature in Section 1.2, after which Sections 2 and 3 examine best of breed approaches to modeling leaders and followers and how they interact in the game framework we have assembled. In Section 4, we summarize some of the games we have authored and analyze the results. Finally, Section 5 discusses what has been learned in striving for the 4 aims, and a research agenda for improving the field of socio-cultural

games and agent based leader-follower modeling and simulation. First, however, we pause to examine the basic socio-culture game of interest here.

A. Socio-Cultural Game Theory

How can an analyst or trainee devise policies that will influence groups for the collective good? And what must a socio-cultural game generator encompass? Figure 1, explained below, attempts to portray a universal class of leader-follower game that groups often find themselves in and that are worthy of simulation studies. This could be for competing groups in a crowd, in an organization, in a region or nation, or even between nations. Analysts would need an appropriate suite of editors and a generator, to help them rapidly model such conflict scenarios and analyze what outcomes arise from different courses of action/policies.

Figure 1 – Overview of the Basic Leader-Follower Game That is Repeated Around the World

Specifically, the socio-cultural game centers on agents who belong to one or more groups and their affinities to the norms, sacred values, and inter-relational practices (e.g., language, gestures, social rituals) of those groups. Specifically, let us suppose there are N groups in the region of interest, where each group has a leader archetype and two follower archetypes (loyalists & fringe members). We will say more about archetypes shortly, and there can certainly be multiple leaders and followers, but we stick in this discussion to the smallest subset that still allows one to consider beliefs and affinities of members and their migration to more or less radical positions. There is an editable list of norms/value systems from which each group's identity is drawn. The range across the base of Figure 1 shows an example of a political spectrum for such a list, but these could just as easily be different parties in a common political system, diverse clans of a tribe, different groups at a crowd event, and so on. Each entry on this list contains a set of properties and conditions that define the group, its practices, and entry/egress stipulations. The authority of the leader in each group is also indicated by a similarly edited list depicted illustratively across the top of Figure 1.

The vast majority of conflicts throughout history ultimately center around the control of resources available to a group and its members. The resources of each group are illustrated along the left side of Figure 1 and are summarized for brevity into three tanks that serve as barometers of the health of that aspect of the group's assets –

(1) economic goods available to the members (jobs, money, foodstuffs, training, healthcare etc.); (2) security and rule of law applied in the group as well as level and type of security available to impose will on other groups; and (3) popularity and support for the leadership as voted by its members. Querying a tank barometer value in a culture game will return current tank level and the history of transactions or flows of resources (in/out), who committed that transaction, when, and why (purpose of transactional event).

Initially, the game is started with manually coded, but empirical based alignments. However, these will evolve dynamically as play unfolds. Specifically, each group leader, in turn, examines the group alignments and notices Loyal Ingroup (A), Resistant Outgroup (C), and those “undecideds” in middle (B) who might be turned into allies. Also, if there are other groups, they are examined to determine how they might be enlisted to help influence or defend against the out-group and whatever alliance it may have formed. Followers’ actions are to support their leader’s choices or to migrate toward another group they believe better serves their personal value system. Actions available to Leader of A are listed in the table on the right side of Figure 1 as either speech acts (motivate, threaten, form pact, brag) or more physical/political acts. Of the latter, there are 6 categories of strategic actions. The middle two tend to be used most heavily by stable, peaceful groups for internal growth and development. The upper two are economic and militaristic enterprises and campaigns taken against other groups, while the lower two categories of actions are defensive ones intended to barricade, block, stymie the inroads of would be attackers. The right hand column of the action table lists examples of specific actions under each of these categories – the exact list will shift depending on whether the game is for a population, organizational, or small group scenario. In any case, these actions require the spending of resources in the tanks, with proceeds going to fill other tanks. Thus the culture game is also a resource allocation problem. Leaders who choose successful policies will remain in power, provide benefits for their followers, and ward off attackers. Analysts and trainees interacting with this game will have similar constraints to their policies and action choices.

B. LeaderSim Game World

The model presented in Figure 1 is the lead author’s theory of role playing socio-cultural games that guides various implementations we have attempted to date [6]. It provides a point of departure for delineating the 5Ps of any scenario – i.e., people (roles, relationships), plot, place, processes (campaigns, actions, speech acts), and player pedagogy. We have authored implementations of this theory that run inside of other people’s videogame engines

such as a Mogadishu Black Hawk Down recreation inside Unreal Engine (vanLent et al. 2004, Silverman et al. 2006b), or an Iraqi food relief scenario inside of the BigWorld Engine (MacDonald et. al, 2006). Such 3-D animated worlds are useful for training of street scenes, however, the graphics become an unnecessary burden when trying to analyze and understand what lies behind such conflicts.

The lead author spent much of 2004 assembling a paper-based version of Figure 1 as a role playing diplomacy game and play-testing it with analysts: Silverman, Rees et al (2005). The goal of the game is to help players to experience what the actual leaders are going through, and thereby to broaden and deepen their understanding, help with idea generation, and sensitize them to nuances of influencing leaders in a given scenario. The mechanics of the game place the player at the center of the action and play involves setting objectives, figuring out campaigns, forming alliances when convenient, and backstabbing when necessary. This is in the genre of the Diplomacy or Risk board games, though unlike Diplomacy, its rapidly reconfigurable to any world conflict scenario.

After completing the mechanics and play-testing, three implementations of the game were created: (1) a software prototype called LeaderSim (or Lsim) that keeps world scenarios and action sets to the simplest possible so that we can easily build and test all of the core ideas of the theory; (2) a scaled up version called Athena's Prism that has been delivered as a fully functioning computer game in mid 2005, though AI opponent features are continually being added; and (3) a streamlined version of the paper-based game has been turned into a boardgame called BigWig©, aimed at being played to conclusion within an hour (it is thus intended to serve as an intro to the diplomatic strategy genre for new players).

The computer implementation of the full game (Athena's Prism) has been developed as a distributed multi-player game architecture. The operations concept, is that this could be played by (i) all human (ii) all artificial intelligence (AI), or (iii) any combination of humans and AI.

In the Lsim game, we consider a world made up of at most 3 abstract territories, each of which has a constituency with up to 3 resources categories. The three lands must be named Redland, Blueland, and Yellowland, though actual leaders and groups can be assigned to these lands through editors. The left side of Figure 2 shows all 3 constituencies (lands) each with 3 sets of resources. Each leader/player gets to own a constituency which may or may not be shared with other leaders in the same territory. Constituencies are marked up according to identity theory [34] with the features and expectations of the leader's constituency (e.g., actions and leaders that the constituency favors or opposes).

In Lsim, resource levels are represented by tokens (bingo chips, shown as histograms in Figure 2) that may be used to pay for various actions one wishes to perform in the world. This resource-action-token representation provides an extremely concise, yet flexible and powerful way of describing scenarios and for carrying out game strategies. Adopting this formalism has allowed us to elicit and represent much of the complexity of situations that are difficult for experts even to verbalize. It is only necessary to get the relative levels of the tokens correct to capture the nature of socio-cultural, resource allocation conflicts between groups. Figure 2 depicts a region filled with two relatively smaller groups (Redland, Blueland) and larger group Yellowland. In some runs, students have set these up to be Richard Lionheart, Guy, and Saladin of the Third Crusade, but it is easy to alter this.

Figure 2 – Lsim: a World Diplomacy Role Playing Game for Humans or Agents

As the base of Figure 2 also shows, there are a set of 3 hostile and 3 non-hostile actions in Lsim. The rules govern the spending (tapping) and replenishment (untapping) of tokens (it's a zero sum game with constant sum of tokens to go around?), how to grow your assets and expand around the globe, how your actions affect other players, ways to defend against attacks, what may be done during communiqués or summits, how espionage works, the carrots or benefits you can offer other leaders, and so on. A graphical action wizard helps step human players through these rules whenever they take a turn. Despite there only being 100 possible actions, the choice set rapidly becomes intractable when one considers the array of possible targets, multiple ways one can pay for any given action, and the various levels one can choose to do an action at. Hence, we restrict our attention in this article to Lsim's limited action set. Also, we omit description of the dialog and spying screens and assets.

In general, when humans play the game, they rapidly evolve a portfolio of strategies that they tend to pursue in parallel, where a strategy is a high level goal that might be implemented by any of a number of alternative actions. An 'action' is defined as a sequence of low level moves governed by the rules of the game. There are only a few moves (e.g., tap/untap tokens, re-assign tokens to resources, etc.). This portfolio or strategy-action-move hierarchies tend to reflect the culture and personality of the leader in a scenario as they attempt to navigate the 'game' against the other players. As one example of this, the middle of Figure 2 depicts a turn where Redland (Richard) is attacked economically by his supposed ally (Guy, the Latin King of Jerusalem). Since outcomes are probabilistic, the game engine handles the 'dice rolling' once both sides have picked a level of attack and defense, respectively, and in the

bottom middle and then the right side of Figure 2, one can see that Guy successfully weakened Richard's campaign and made him more dependent on his host (Guy), though certainly not more sanguine about the alliance.

For the AI to be able to replace a human player and to assemble and manage a portfolio in a way as to reasonably emulate a world leader, a number of components are required in the mind of the agent as shown as the next few subsections amplify. In particular, Performance Moderator Function Server (PMFserv) is a human behavior modeling framework that manages an agent's perceptions, stress and coping style, personality and culture, social relationships, and emotional reactions and affective reasoning about the world: Silverman et al.(2002a, 2002b, 2005, 2006a, 2006b).

C. Computational Theories of Agents

To our knowledge there are no other agent based (or other) approaches that come close to embracing the breadth and depth in Figure 1. Our approach is based on complex, descriptive agents based on incorporating human behavior models and other social systems models. As one approach, game theory itself offers some useful notions, though for the most part it is a rich theory only for simple games and is incapable of handling the details of Figure 1. Given that the computational complexity of coalitional games [typically $O(NP\text{-complete})$ or "non-deterministic polynomial time"] preclude analytical solutions (Woolridge & Dunne, 2004), for formalisms such as partially observable Markov decision processes (POMDPs), the game is computationally intractable except via approximations. Behaviorally inspired, descriptive models, on the other hand, perform better. A comparison study between approximations of prescriptive approaches vs. "descriptive approaches" (Simari & Parsons, 2004) shows that as games grow larger and more complex, the descriptive approaches outperform prescriptive counterparts and tend to provide closer convergence and faster performance. Any models built in formal paradigms tend to limit themselves to stylized, simple problems, such as 2 player games with minimal action choices.

As a second approach, AI in existing commercial videogames is devoted to the use of finite state machines for low-level functions such as navigating, path finding, and personal combat tactics. A small portion of this community focuses on what might be called "Leader AI" mostly for strategy games, but also role playing games and a few massive multiplayer online games. Most often these games are based on fantasy or futuristic alien world, where realism is not important and never achieved. The leader AIs tend to be the opponent leaders and they fairly uniformly rely on simple rules buried in scripts to govern the branching of agent behaviors, with the AI employed to entertain the player. In these games, rules are kept simple and pre-scripted, often with an easy to beat opponent for

new players and a harder level to challenge better players (accomplished by making the AI omniscient). The bottom line for the videogame community is that it does not pay to invest in leader or follower AI beyond the minimum needed to engage the player (buyer), and there are many instances where this simple approach suffices to create enjoyable play and sufficient drama for entertainment purposes, keeping players engaged for many hours. However, if one is interested in player learning (as we are here) in more realistic, strategic situations, one cannot ignore the need for AI's behavioral realism computational performance.

All the issues mentioned above lead us to the need to have our own game generator – composability for rapid mock up of real world socio-cultural conflict scenarios, playing against realistic leader-follower agents, and flexibility to use for analytical purposes. This does not preclude use of existing game environments, and as will be mentioned at key points in this article, we have plugged our agent-based human behavior simulator (PMFserv) in to drive agents in other game worlds – e.g., Unreal Engine, BigWorld, JSAF, etc.

An important learning objective of our research program is to directly model and explore the underlying behaviors and motivations of leaders and followers. For these purposes, we turn to a final agent sub-community -- that of the human behavior modelers. Unlike game AI, human modeling need not focus on navigating and battling around spaces, but instead invests in detailed modeling and simulation of individual agent's decision and cognitive processes, and worries about accuracy and validity of the resulting behaviors. This community eschews the rationalistic-normative paradigm of game theory, and instead often embraces an array of descriptive, naturalistic, or heuristic approaches to modeling human cognition and reasoning: e.g., see Pew & Mavor (1998). As cited above, descriptive/cognitive approaches: (1) hold a better chance of converging in coalition games; (2) are more tractable for plan recognition/intentionality modeling; (3) offer improved mechanisms for the perception and reasoning about aspects of other agent behavior such as trust, relationships, and group membership choices; and (4) hold the promise to help agents move beyond surface features and explain their choices and behaviors. Despite these arguments for richer human behavior models, the human behavior modeling community does not hold a well-developed approach to all these issues. There are some widely used cognitive models (e.g., SOAR, ACT-R) that tend to be quite adept at logical processes of reasoning and learning but which ignore many other components of behavior and cognition pertaining to individual (values, personality, physiological and cognitive constraints) and groups (affinity, relationships, social cognition, culture). In contrast, PMFserv complements such models of cognition with a wide range of implemented best-of-breed human Performance Moderator Functions (PMFs) such as: for physiologic

impacts (e.g., exertion, fatigue, injuries, stimulants, noise); for sources of stress like event stress or time pressure; and with models of emotion and affect to modulate expected utility formulations and to help account for cultural, personality, and motivational factors that impact upon leader-follower relations, Silverman (2005). PMFserv also manages social relationship parameters and thus macro-behavior (e.g., in collectives or crowds of PMFserv agents) emerges from individuals' interactions and micro-decisions. As a result, PMFserv is the agent framework adopted for our leader-follower AI. The human behavior modeling community is relatively new to the idea of modeling individual differences drawn from validated personality/culture models and instruments. That is the portion of PMFserv we focus on here, as the next two sections explain for leaders and followers, respectively.

II. THEORIES ON THE PERSONALITY AND VALUES OF LEADERS

An important source of ideas for the modeling of leaders comes from the field of political psychology. While these are not implemented or operationalized agent models, they can be a source for creating agent frameworks with greater degrees of realism. In terms of leader theories that might impact the leader-follower game (Aim 2 of studying leader –follower models), Chemers (1997) indicates that while the leadership literature often appears fragmented and contradictory, it can be summarized in terms of four dominant theories of leadership: Leader Trait Theory, Transaction Theory, Transformational Leader Theory, and Situational Contingency Theory. These leadership theories seek to stipulate the most effective leadership style and behaviors that a leader must perform well in order to be effective in a given situation. Our current Lsim agents include implementations of some portions of the situational, transformational and transactional features as will be outlined shortly. However, these theories ignore leader to leader relationships, plus they tend to be normative and prescriptive. From our viewpoint, we need to embellish these with a theory that focuses upon inter-leader dynamics and that permits us to model leaders as they are (not as a prescription would prefer). As a result, we turn now to a descriptive theory of leader style, one that is measurable and can be fully implemented in our agent-based framework.

After two decades of studying over 122 national leaders including presidents, prime minister, kings, and dictators, Hermann (1999) has uncovered a set of leadership styles that appear to influence how leaders interact with constituents, advisers, or other leaders. Knowledge about how leaders react to constraints, process information, and are motivated to deal with their political environment provides us with data on their leadership style. Hermann determined that the traits in Table 1 are particularly useful in assessing leadership style.

In Hermann's profiling method, each trait is assessed through content analysis of leaders' interview responses as well as or other secondary sources of information. Hermann's research also has developed methods to assess leadership at a distance, based mostly on the public statements of leaders. While both prepared speeches and statements from interviews are considered, the latter is given preference for its spontaneity. The data are collected from diverse sources, usually as many as 50 interviews, analyzed or content coded, and then a profile can be developed. These are then compared with the baseline scores developed for the database of leader scores. Hermann (1999) has developed mean scores on each of the seven traits. A leader is considered to have high score on a trait, if he or she is one standard deviation above the average score for all leaders on that trait.

Table 1 – The Seven Traits of the Hermann Leadership Style Profile

Our Lsim implementation includes a full set of screens and sliders for quickly tuning these parameters for each leader in a conflict scenario. Due to space limits we omit presenting these here, but interested readers should consult Silverman et al. (2005). Using our Hermann implementation moves us toward Aim 2 (models of leader – follower behavior) as one can populate a game with real leader profiles provided the profiling were done properly. Thus, for example, one could determine which leaders tend to be deceitful vs. honest. Specifically, the leader with low belief in control (trait 1) but high need for power (trait 2) tends toward deceit, while the leader with high trait 1 and high trait 2 tends toward accountability and high credibility. Likewise, the same could be done for the other traits (and our new trait of protocol vs. substance), as we will attempt to demonstrate.

Any implementation of a best of breed, paper-based model runs into the “catch-22” (fragmented, few computationally useful models in social sciences, yet it is the computational implementation that would integrate and make social system models useful) outlined in Aim 3 (improving the science by iteratively understanding, synthesizing and testing relevant best of breed social science models), and the need to try and enhance that model by integrating it with others and making it executable by autonomous agents. Hermann gives no guidance on how an agent feels if it is achieving its need for power (trait 2). We turn in the next section to a discussion of some of the mathematics needed to complete the model and attempt to interpret it faithfully, yet provide a software implementation. Subsequent sections examine our interpretation in terms of face validity, Turing tests, correspondence tests, and sensitivity analyses.

A. Agent Personality, Emotions, Culture, and Reactions

In Lsim, each leader is modeled within a framework known as PMFserv (Silverman 2005) where the leader's cultural values and personality traits represented through a Goals, Standards and Preferences (GSP) tree shown in Figure 3. These are multi-attribute value structures where each tree node is weighted with Bayesian importance weights. A Preference Tree is one's long term desires for world situations and relations (e.g., no weapons of mass destruction, stop global warming, etc.) that may or may not be achieved in the scope of a scenario. In Lsim agents this translates into a weighted hierarchy of territories and constituencies (e.g., no tokens of leader X in resource Y of territory Z). When faced with complex decision spaces, different individuals will pursue different long-term strategies which, mathematically, would be very difficult to compare objectively. Chess players, athletes, and scientists develop their own styles for solving the types of problems they encounter. We make use of the *preference* structure of an agent to account for much of this. For example, one can say that a particular chess player *likes* or is comfortable with certain configurations of the pieces on the board. This allows for the expression of long-term strategic choices that are simply a question of style or preference as to how the world should be.

The Standards Tree defines the methods a leader is willing to take to attain his/her preferences. Following from the previous section of this article, the Standard tree nodes are mostly Hermann traits governing personal and cultural norms, plus the additions of Protocol vs. Substance, and top level guidelines related to Economic and Military Doctrine. Also, we add two standards from the GLOBE study (House, 2004) on 'Scope of doing Good' and 'Sensitivity to Life (humanitarianism)' to the Standard tree. 'Scope of doing Good' defines the extent to which the leader is self-serving and ensuring the safety and security of the individual or inner-circle at the expense of the group he or she leads. 'Sensitivity to Life' (Humanitarianism) refers to the humane orientation trait characterized by considerate behavior that includes compassion and generosity.

Personal, cultural, and social conventions render inappropriate the purely Machiavellian action choices ("One shouldn't destroy a weak ally simply because they are currently useless"). It is within these sets of guidelines that AI could be improved significantly to provide realistic behavior and solutions. Standards (and preferences) allow for the expression of strategic mindsets. When a mother tells her son that he shouldn't hit people, he may not see the immediate tactical payoff of obeying. However, this bit of maternal wisdom exists and has been passed down as a standard for behavior precisely because it is a non-intuitive strategic choice whose payoff tends to derive from

what *doesn't* happen far into the future as a result. Thus, our framework allows our agents to be saved from their shortsighted instincts in much the same way as humans often are.

Figure 3 – GSP Tree Structure, Weights and Emotional Activations for Country Leader

Finally, the Goal Tree covers short-term needs and motivations that implement progress toward preferences. In the Machiavellian and Hermann-profiled world of leaders, the goal tree reduces to a duality of growing vs. protecting the resources in one's constituency. Expressing goals in terms of power and vulnerability provide a high-fidelity means of evaluating the short-term consequences of actions.

With GSP Trees thus structured, we believe it is possible to Bayesian weight them so that they will reflect the portfolio and strategy choices that a given leader will tend to find attractive, a topic we return to in Section 4 of this paper. As a precursor to that demonstration and to further illustrate how GSP trees represent the modified Hermann profiles, consider the right side of Figure 3. There we see the weighted GSP tree of Saladin. Other papers discuss how the weights may be derived so as to increase credibility: e.g., see Bharathy (2006), Silverman (2002a,b, 2006 pt.2). Here it is more pertinent to discuss how the G-tree implements the Hermann power vs. protect trait. Beneath each subnode that has a + sign, there are further subnodes, but under the G-tree (and P-tree) these are just long sets of constituency resources with importance valuated weights and hence they aren't show here. The standards or S-tree holds most of the other Hermann traits and their important combinations, such as traits 1 and 2 that combine to make the four subnodes covering all possibilities of Belief in Control vs. Need for Power. Likewise, there are subnodes for In-Group Bias (or Treatment of Outgroups), Degree of Distrust, and. Openness, as mentioned earlier, is a direct replacement for two other traits, while Task vs. Relationship focus is also supported. The modifications to Hermann show up as the Protocol vs. Substance subnodes and the key resource specific doctrines of importance to that leader. In Saladin's case most weights are fairly balanced, while in Richard's case, the G-tree weights show he leans heavily toward power and growth which is also consistent with his P-tree weights on his own resources. Richard's standards reveal him to be Low Openness, Substance- and Task-focused, and favoring asymmetric or non-conventional attacks (he did slaughter thousands of unarmed townsfolk).

Just to the left of the weight value on each node of the GSP trees of Figure 3 are two "reservoirs" that reflect the current activation of success and failure of this node, respectively. These reservoirs are activated and filled by events

and states of the game world as observed by the agent. Figure 3 shows the start of a game where Saladin has yet to be attacked and he perceives the world in a fairly satisfied light. In general, we propose that any of a number of k diverse activations could arise with intensity, ξ , and that this intensity would be somehow correlated to importance of one's GSP values or node set (GSP) and whether those concerns succeed or fail for the state in question. We express this as

$$\xi_k(\mathbf{b} \in \mathbf{B}) = \sum_{j \in J_k} \sum_{v \in V} [W_{ij}(v \in V) * \Phi(r_j) * \zeta(v) * \psi] \quad [1]$$

Where,

- $\xi_k \rightarrow \xi_k(\mathbf{b} \in \mathbf{B})$ = Intensity of activation, k , due to the b th state of the "B" world.
- J_k = The set of all agents and objects relevant to k . J_1 is the set consisting only of the self, and J_2 is the set consisting of everyone but the self, and J is the union of J_1 and J_2 .
- $W(v \in V)$ = Weighted importance of value set V to the agent.
- V = The set of Goal, Standard, and Preference held by the agent.
- $\phi(r_j)$ = A function that captures the strength of positive and negative relationships one has with agent or object j that are effected or spared in state b .
- $\zeta(v)$ = degree of activation for a Goal, Standard, or Preference
- ψ = A function that captures temporal factors of the state and how to discount (decay) and merge one's GSP activations from the past (history vector), in the present, and for the future

It is important to note that the weights adhere to principles of probability; e.g., all child node insights add to unity beneath a given parent, activations and weights are multiplied up a branch, and no child has multiple parents (independence). Although we use fixed weights on the GSP trees, the reservoirs serve to render them dynamic and adaptive to the agent's current needs. Thus, when a given success reservoir is filled (that concern having been satiated in the immediate past), that tends to nullify the importance of the weight on that node (or amplify it if the

failure reservoir is filled). In this fashion, one can think of a form of spreading activation (and deactivation) across the GSP structure as the scenario proceeds.

According to other best-of-breed models (Damasio, 1994; Ortony, Clore, and Collins, 1988; Lazarus, 1991), our emotions are arousals on a set of values (modeled as trees) activated by situational stimuli as well as any internally-recalled stimuli – e.g., see full descriptions of these models in Silverman et al. (2002a,b, 2006a,b). These stimuli and their effects act as releasers of alternative emotional construals and intensity levels, and they assist the agent in recognizing problems, potential decisions, and actions. According to the theory, the activations may variously be thought of as emotions or subjective (moralistic) utility values, the difference being a matter of semantic labeling. Within such a framework, simply by authoring alternative value trees, one should be able to capture the behavior of alternative “types” of people and organizations and predict how differently they might assess the same events, actions, and artifacts in the world around them.

B. Agent Decision Making

What is missing in the previous section is how an agent notices the game world, moves of others, and sense of situation. This discussion will illustrate how this happens using one Hermann factor (power and Vulnerability) as an example. Hopefully, it is fairly straightforward for the reader to extend that to how the other factors are also deployed. Full details exist in Johns (2006).

Central to a given leader’s G-Tree reasoning is its perceptions of who threatens it and/or whom it’s vulnerable to. Likewise a given leader may be equally interested to estimate who can it influence to best increase its resource assets and thereby its power in the world. Obviously, GSP tree weights will govern how aggressively a given leader pursues each of these Vulnerability vs. Power concerns, however, we assume that all leader agents need to be able to compute how vulnerable and/or powerful they are at each turn of a game by inferring or assuming other agents’ vulnerabilities. Since the game rules define precisely which resources can be used to take hostile actions against which other resources, one can derive a measure of a player’s *Vulnerability* directly from the state of the game world and the rule set. A leader can compute utility derived by other agents and build world utility information for various key alternatives by assuming or inferring (not necessarily accurately) sets of GSP tree weights. Intuitively, by factoring Vulnerability into these world utility calculations, an agent can avoid world configurations in which

another is poised to conduct a devastating attack. Adding border defenses, stocking up on supplies, and pulling money out of the economy can all be viewed as behaviors motivated primarily by Vulnerability management. The Vulnerability formula (β) works by generating the percentage of a given player's tokens that can be expected to be lost to a given player in the coming round of attack actions (a_i). For each hostile action ($a_i \in A$) that can be initiated by another player (g), the number of tokens available to attack and defend is tallied. From this the probability of victory is determined, and then multiplied by the percentage of tokens vulnerable to this attack versus the total number owned by the vulnerable player in each resource category. This is the expected percentage of tokens to be lost if this attack occurs in the next round. The maximum over all attacks, then, gives this player ℓ 's

$$\text{Vulnerability score } \beta \text{ to player } y. \beta_{xy} = \text{Max } a \in A \left\langle \text{Pr}(a) * \frac{\sigma(x, a)}{C(x)} \right\rangle \quad [2]$$

Agents who purely manage Vulnerability, while interesting in their behavior, are not entirely realistic. Human players tend to balance Vulnerability against its inverse, *power*. Where Vulnerability measures the expected number of tokens a player can lose to other players in the coming round, power measures the expected number of tokens a player can take from others. The calculation of the power heuristic is exactly the opposite as for Vulnerability. Player A's Vulnerability to Player B is the same as Player B's power over Player A.

Taking the leader's perceived difference between Power and Vulnerability provides a surrogate for the leader's overall sense of utility ($U(G_x)$) of the current state of the world, G , when divorced from his value system and other factors: $U(G_x) = \alpha_x - \beta_x$ [3]

Recall, however, that a given leader agent (1) tracks who is aligned with whom, tallying things like trust, (2) monitors all resource levels and who used what actions upon them, and (3) its own actions to achieve its long term preferences or P-tree, as modulated by its standards. The α and β , thus calculated, serve primarily as activations on the leaf nodes of some of the GSP tree branches. PMFserv uses a wide assortment of similar activation mechanics for other factors and computes the Expected Utility (EU) of the world and of new action possibilities when projecting next steps. That is, PMFserv serves as the point where diverse GSP personality and cultural value sets, stressors, coping style, memories, and perceptions are all integrated into a decision for action (or inaction) to transition to a new state (or remain in the same state) and to determine the portfolio of strategies-moves-actions that best maximize that agent's GSP Tree values as follows:

$$\text{Max EU}(a) = \sum_{b \in B_a} UI(b) * pr(b) \quad a \in A \quad [4]$$

Where,

A = action set available after GSP and stress-constrained perception

$p_i(b)$ = probability of action a leading to state b

$$u_l(b) = \frac{\sum_{k \in K} \xi_k(b)}{11} \quad [5]$$

Utilities for next actions, a_k , are derived from the activations on the GSP trees in the usual manner as in Silverman, Johns, et al. (2002a, b) and as Silverman et al (2002a, b) and as briefly summarized for power and Vulnerability here. That is, utility is the simple summation of all positive and negative activations for an action leading to a state. Since there are 11 pairs of oppositely valenced activations in PMFserv's emotion model (Ortony, Clore, et.al., 1988), we normalize the sum so that utility varies between -1 and +1.

C. Intention Modeling and Mirroring Biases

In order to assess speech acts and motivations of others, agents require certain information about the other players of the game. Specifically, an agent needs to be able to evaluate an action from the perspective of another. This is nested intentionality – or – model of mind. For it to work in Lsim, each leader agent seeks to model the goals, standards, and preferences of other leaders to the best of its ability, using starting assumptions (mirroring) updated by observation. It does this so that its subjective expected utility algorithm can determine the likely next moves of the other players it is most threatened by and/or most hoping to influence.

As mentioned earlier, our Aim #2 is to try and implement realistic heuristics in the models. As a starting set of heuristics, Heuer (1999) indicates people commonly use mirroring and confirmation. We accomplish this by modeling the utility functions of other agents as derived from one's own goals, standards, and preferences. Models are initialized through a “mirroring” process, where agents assume that the motivations of other agents are similar to theirs, but inverted. For instance, if L_x wants to destroy L_y 's armed forces, L_x assumes that L_y wants to destroy L_x 's

armed forces as well. These assumptions are updated as action unfolds via observation as follows. Suppose leader L_x performs an action A_y :

- For each resource C_i , determine whether A_y has increased or decreased its Vulnerability β
- If C_i is less vulnerable, he must want to protect it
- If C_i is more vulnerable:
 - If L_x owns C_i , he must not value its protection
 - If L_x does not own C_i , he must want to destroy it

III. MODELING FOLLOWER VALUE SYSTEMS: THE EVOLUTION OF DANGEROUS IDEAS

We introduce three refinements in order to also be able to model the values and motivations of followers – (1) additions to the GSP trees, (2) a group-affinity profiling instrument, and (3) group transfer dynamics ((3) has been left out of the scope of this paper). In keeping with Aim 2, each of these refinements is an implementation of a well-respected model drawn from the social sciences (Hofstede,2003; House et.al., 2004;Hermann, 1999). In this section, we describe this integration and synthesis in terms of an example individual Villager from a South East Asian country. He grew up as a Muslim in a society with a Buddhist majority (BlueLand on right of Figure 4), where the Buddhists occupy the cities, hold the elected posts, and get the best jobs. Most of the Muslim's are villagers in the rural provinces, a region they would like to separate and have autonomy over (shown as Yellowland), The elder villagers tend to be more moderate and compliant, while the younger ones, like the one modeled in Figure 4, were sent to Wahhabi schools elsewhere, went to college for training, and have returned to find there are no sources of employment that will use their college training (Yellow's economy is very low in Figure 4). Further, their religious practices are frowned upon, their community is unable to open Wahhabi schools officially (these are being operated by the returned college grads in their homes), and their initially peaceful protest events have been met with police brutality (Blue ArmedForces present in Yellowland). The bottom left box of Figure 4 shows the country leader has been discriminating against Muslim villagers, the moderates disagree with these policies, while this individual is opting to "oppose". He is just at the juncture of having switched groups from moderate villagers to the local insurgency (Red tokens on right side in YellowLand) which itself is on the verge of connecting with radical extremist groups, if the discrimination continues. Let us now see how the three sets of refinements allow the PMFserv model to capture such a radicalization.

In terms of the GSP tree changes, as the right side of Figure 4 reveals, there are additional nodes on the Goal and Standards trees. While Figure 4 does not show the full trees, the previous goals and standards are still there but may be zeroed out if this is strictly a follower, or may be left in at some degree of importance if this is a mid-level leader. For villagers in general, where day-to-day existence is a struggle, Maslow's hierarchy of needs is considered a useful representation of the range of short term goals that a person might have to be concerned about. Each of these nodes are activated by lower level branches on the tree, tanks in the physiology and stress model, or the affinity instrument described below. Clearly all of this individual's basic needs are failing to be met, and the Maslovian nodes are all activated with negative stimuli in this example, accounting for the extreme distress level of the goal emotions on the left of Figure 4. In terms of the Standards Tree, the Hermann factors are all still there (but for example, the leader specific goals are set negligible weights), but two new nodes on the tree exist that capture several Hofstede (2003) and House (2004)'s GLOBE study cultural factors. Conformity Assertiveness is a way to capture Hofstede's Power-Distance and Individualism factors (respect authority, conform to society) with the GLOBE study's Assertiveness factor. Likewise, the Humanitarian node is a GLOBE factor. The weights on Figure 4 for this villager's standards reflect his idealism, training in religious and college settings, and his alienation. Hence he differs from the moderate villagers in the various Hofstede and GLOBE factors, and also his personal 'military doctrine' does not shun violence. According to Sageman (2004), these are the seeds, but they do not alone explain performing violent and suicidal acts. The activations on his Standards tree tanks must also be such to tip him from just disagreement and protest over to actual physical opposition and violence. We turn now to the second set of refinements that enable radicalization and provide the needed extra activations for these tanks.

Figure 4 – Screens from Lsim and PMFserv Depicting a Buddhist-Muslim Conflict in SE Asia

Specifically, for determining an individual's group affinity, one needs an instrument that measures it. The instrument that we have adopted here involves Eidelson and Eidelson (2003) who have developed a five belief ("dangerous ideas") framework for better understanding the psychology of individual-group dynamics particularly relevant to conflict settings. These five beliefs revolve around issues of Helplessness, Vulnerability, Injustice, Distrust, and Superiority (in short H|VID|S):

- *Helplessness (H)*: Refers to perceived inability to influence or control events and outcomes; self-perpetuating because it diminishes motivation.
- *Vulnerability (V)*: Revolves around a sense of living in harm's way amid constant threat and peril.
- *Injustice (I)*: Reflects perception of being victim of mistreatment by specific others or by the world at large.
- *Distrust (D)*: Focuses on the presumed hostility and malicious intent of other individuals or other groups.
- *Superiority (S)*: Revolves around conviction of being better than others—morally superior, chosen, entitled, or destined for greatness.

Eidelson has developed the Individual-Group Belief Inventory (IGBI) to measure these five beliefs at three different levels of analysis: (1) an individual's personal beliefs about his personal world; (2) an individual's personal beliefs about a specific group in which he holds membership; and (3) an individual's personal perceptions about the shared beliefs this group holds about itself. In the simulation work described here, we have focused on the second of these three levels, namely the beliefs that individual group members hold about their group (in future work we may attempt to more fully implement the other levels of analysis as well). These beliefs are considered particularly important influences on a group member's perceptions of his/her group's current circumstances and future prospects. In our current modeling work, we have simplified the dangerous ideas framework into three domains by combining the Vulnerability, Injustice, and Distrust domains into one broader domain (hereafter labeled VID) representing a member's belief that his/her group has legitimate grievances against a threatening adversary.

The PMFserv implementation of this model lies in marking up diverse leaders and groups with all possible levels of "dangerous ideas" (DI) that they could afford to different potential members and perceivers of their group. Depending on the perceiver, two individuals may view the same group as superior or inferior, as suffering grave injustices or as exaggerating minor slights, as helpless or capable of effective action, and so on. The markups expose all the possibilities and include universal rules that apply to all viewers. Thus if a viewer sees a group as helpless and he doesn't want to be helpless, there will be negative activations for remain loyal to this group and following its action choice policies. Conversely less negative (and possibly positive) activations will be afforded for the member who exits away from this group perceived as helpless. The mechanics of how PMFserv handles perception have been covered in Cornwell et al. (2003), Silverman et al. (2005, 2006a) and we won't go into them in full detail here except to mention that once markup is done, perceptual types are driven by the agent's emotional and utility construals of the current state of the world. For example, when Standards and Preferences are perceived to be

violated Injustice and the anger emotion are aroused, while Vulnerability comes from the Safety failings on the GSP tree, and emotions about that. Part of a Muslim villager's perception of the country leader is as shown in the PMFserv screen of Figure 5 which reflects a range of VID parameter settings (neutral, very low, all the way to High), and the fact that the villager sees the leader as causing VID High. This perception on the part of the villager is based on observing actions that the leader has taken to discriminate against Muslims (no jobs), to deny their request for Wahhabi schools, to brutally put down protest, etc. Conversely, Buddhists who view this leader see him as neutral or low VID. The former would be for those in his political party, while the latter would be others who pursue normal political processes to resist his positions. Likewise, the Muslim villagers and Buddhists will form similar judgments about the various Muslim groups/leaders in the scenario.

Figure 5 – Example of a Villager's Range of DI Perception Possibilities When Viewing a Given Leader

Up to this point we have seen how an agent's GSPs and perception of the range of possible DI markups permit him to examine the state of the world and assess the utility and emotional construals that each group/leader affords up to that point. As VID rises, members may express their support or opposition to the leader on a scale of perceived grievance (GR). The GR_{ij} is the expressed grievance level by agent i relative to group j . This is computed as the negative emotions activated when agent i perceives the VID levels of a given group or leader's actions. GR is thus feelings about the world due to actions that violate one's GSPs and that cause negative emotional utility. In the current model, expressed Grievance is fed from the levels of the goods and security tanks and the individual actions of the leader to affect what these contain. Of course, perceived Grievance alone does not make a rebellion, but it might contribute towards the rebellion. The Grievance should be combined with a feeling that the other group or in this case, the leader or the central authority, induces helplessness, vulnerability, injustice, distrust, and/ or attack on superiority. As the leaders did with Figure 1, the followers similarly take each set of opposing groups and place them along a scale as shown below. The decision that the villagers make is expressed as grievance, where the grievance is in the scale of -4 to +4 are given below (also shown are the Grievance State IDs of the simulation of Sect 4.2):

Table 2 – The Grievance Scale Embedded in Follower Agents

In classifying the actions into the abstract scale, the emotive and non-emotive utility associated with the action was assessed to be the criteria for grouping actions together. The actions the villagers would perform along with the grievance values on an abstract scale are given below. These values (abstract scale) indicate the level of support or opposition to the government.

These actions are on abstract scale, which ranges from total support of the majority that is oppressing you (if you can't lick them, join them), to being undecided and/or helpless in the middle, to the other extreme of supporting and ultimately joining the insurgency. On the left the extremes on either end, the agent will submit to militaristic commands of the leader of that group, while at the next level two lower levels they will be only willing to go to protests, and verbally and economically support the activities of that group's leaders. Thus, every state all the way through GS4 represents Voice.

In terms of our Villager from earlier Figure 4, he has just transitioned from the undecided or neutral Muslim villagers category to the OPPOSE level (GS3). He saw Country Leader as creating high VID.

IV. SOCIO-CULTURAL GAME RESULTS TO DATE: TURING AND CORRESPONDENCE TESTING

The previous section delineated the human behavior model of leaders and followers in terms of value trees and activation mechanics in PMFserv. Here we turn to an examination of how these archetypes work in a game. The results about to be presented serve several purposes. First they are part of a test as to whether we can implement personality profiling frameworks such as Hermann and DI with any success (Aims 1 & 2), and whether our extensions are effective (Aim 3). Also, these results provide insight into what is needed as next steps. Next, these results offer validity assessment insights of our cross-cultural Sims including Turing and historical correspondence tests. Finally, once validity is established, it is interesting to use the models and gameworlds to explore what-ifs and sensitivity of outcomes to alternative competing action policies, leader styles, and other group parametrics (Aim 4).

A. Lsim: Fighting Due to Misconceptions, Exploring Alternatives

In this section we configure the game in Lsim to involve 3 autonomous leader agents - Yellow (Saladin), Red (Richard Lionheart) and Blue (Guy of Lusignan), each with 3 resources as shown in Figure 2. There are no human

role players, and followers are just tokens under the control of the leaders. The GSP trees of the leaders are as already described.

In this section, the longitudinal plots are strip charts summing the strength of all 3 resources over time (turns are a surrogate for time) for each leader. Being stochastic systems, repeating the runs of the game with identical starting conditions such as above can produce different results due to chance and randomness of action outcomes. The outcomes (I & II) in Figures 6 are example runs illustrating two possible futures from the similar starting conditions.

1. Correspondence Tests -- The actual course of the Third Crusade involved most battles going in Saladin's favor, and in the end Richard only secured the rights of Christians to live in and visit Jerusalem. Guy lost his right to the throne of the Latin Kingdom of Jerusalem when his wife died, and was banished to Cyprus. Figure 6a seems remarkably in accord with that history. It begins with Yellow (Saladin) being attacked by Red and Blue, but unfolds in Yellow's favor. In the end, Blue (Guy) is eliminated from the game board and Red (Richard) only survives by removing his military and economy and specializing in a non-threatening resource (People).

Figure 6A: Resources vs. Turns for Outcome I

Figure 6B: Resources vs. Turns for Outcome II

However, we are not running a warfare model, but rely instead on dice rolls for the outcome of each individual battle. If we do not bias the battle outcomes to fall largely in Yellow's favor (or add substantially more Yellow military), we can wind up with very different outcomes. Figure 6B shows Red growing more powerful, and all three co-existing by the end, where a somewhat stable equilibrium takes hold. In fact, this seems like the outcome of the First Crusade where the Crusaders did dominate until Saladin amassed a large army, defeated Guy, and retook Jerusalem in the Second Crusade.

2. Turing Test – This test is whether a knowledgeable reader would view the agent's decisions as similar to what the real human would tend to do, and feel, under those circumstances. In a separate paper, we expose the Richard leader to a stimuli – that of the weakly-defended, wealthy city (leader) of Acre. He proceeds to siege the town, demand ransom for sparing it, and finally slaughters all inhabitants when ransom payments are too slow: Silverman & Bharathy (2005). In the current paper, rather than repeat those results, Figures 7A-D illustrate (using Outcome of Fig. 6A) the type of parameters that are tracked and available for inspection. They respectively include (but are not

limited to) Utility, Error of Impression of the Opponents' Preference Tree, Power (or its inverse, Vulnerability), and Preference Achievement. Each of these figures includes an explanation of the 3 leaders. Upon showing these to reviewers, the response is generally favorable. At any juncture, one also can drill into the heads of any of the leaders and inspect the GSP trees and succeed/fail buckets (as in earlier Figure 3)

Figure 7A: Utility vs. Time

Figure 7B: Impression Error vs. Time

Figure 7C: Power vs. Time

Figure 7D: Preference Failure vs. Time

Turing Test and Cooperate-Defect Behaviors –Hobbes (1947) posed the influencing issue with which we opened this article as the central dilemma of political science -- how to persuade individuals to act in the interest of the collective. Too often, this has come to be framed as the Prisoner's Dilemma (PD) game where the implication is that rational, non-trusting individuals will act in their own self-interests (defect) rather than for the collective good (cooperate), even though the latter would improve their utility as well: e.g., see Axelrod (1993). While game theoretic PD formulations are commonly researched and taught, there is recent evidence that this is not necessarily how real people act. In fact, there is now strong laboratory evidence documenting humans' frequent willingness to cooperate even in one-shot PDs where the incentive to defect is substantial and unambiguous: e.g., see Field (2001), among others. Conversely, individuals in tribes (e.g. in the Crusade factions, or in today's Middle Eastern region) are renowned for harboring grudges against each other for years, yet these fade when they are faced by a common enemy (Fandy, 1994).

Such behaviors seem to be spontaneously emerging in the game results of Figures 6-8. For example, due to their perceived preference compatibility, Red and Blue discount their vulnerability to one another when making decisions, and amplify the importance of avoiding vulnerability to Yellow due to the fact that some of their preferences conflict directly with what they perceive to be Yellow's preferences. This creates an implicit pact between Blue and Red, with no acts of communication uttered, in which they will sometimes, in their pursuit of Yellow's destruction, create opportunities for one other to cause harm, and trust that these will not be exploited. For instance, suppose all three leaders have equal strength in military resources and it is Blue's turn to act. Blue may consider a full-scale attack against Yellow, which would ordinarily be hampered by the fact that allocating all of his military resources to such an attack would allow Red to destroy these exhausted forces at will on his next turn.

However, since Blue believes their goals to be compatible, and destroying the resources of those who will use them to advance one's own preferences is unwise, Blue can discount, if not entirely dismiss, this concern. Of course, Blue's model of Red's preferences is imperfect, and consequently this unspoken agreement may be broken when Blue does not expect it. Indeed, in Figure 8A explained below, when Yellow is countered, we see Red and Blue facing off militarily.

B. Impact of Leader Policies on the Evolution of Dangerous Ideas

The scenario in the previous section showed that two groups with congruent GSP values came to believe (mistakenly) that they were threatened by a larger group whose values they misunderstood. There the main concern for YellowLand was to find out how it could live peaceably with its 'opponents'. In this Section, by contrast, we return to the scenario of Sect.3.2 where a larger group of one religion and its leader (Blue) actually do discriminate against the two smaller groups (the moderate Villagers, Yellow, and the rebels, Red) of a different religion, both in terms of values as well as overt policies. The main policy concern here will be to find out: how should Blue leader address this problem so as to attain his own goals yet also be more tolerant and prevent a full blown insurgency from being spawned? Why is violence rising? What are the consequences for domestic politics? What would be the best targets and times to intervene? Unlike the ancient history of the Crusades, this case will draw training and test data from events during the past three years in South East Asia. Only publicly available data sources have been used, but the sponsor requested this land remain unnamed, and hence we will refer to them solely as Blue, Yellow, and Red (see Figure 8A).

Figure 8 – Start and End States During the Correspondence Test: Lsim Summary View

As one can observe, the Muslim population at start is neutral with few grievances registering, whereas towards the end, the population reflects radicalization and spread of nonviolent and violent protest.

Figure 8C – Civil Violence View of Population Membership Before and After Correspondence Test

Figure 9 – Violence Events on Grievance Scale

1. Correspondence Test

There were three types of data/ empirical information employed in this model:

- Numerical data as well as empirical materials on BlueLand, particularly the violent incidents occurring in the rural Yellowland provinces under the control of Blue,
- Empirical information about the decisions made, along with the contexts of these decisions, by the specific personnel being modeled (the Leader of the Blue), and
- Culture specific information for the Blue and Yellow from such studies as GLOBE (House et.al., 2005), as well as religious doctrines affecting the people of concern.

During the 1990s, the country was relatively stable, however, in the last few years, the rural provinces (YellowLand) have seen a rise of Muslim anger against the central Blue government, and the internal security situation in these provinces has rapidly decayed. Certain factions in YellowLand are seeking independence from BlueLand. During 2004, a small group of people, indicated as Red in Figure 8A has committed an increasing number of violent acts against Buddhists (Blue people). The level and sophistication of the attacks has been increasing to the point where people are questioning whether there may be outsiders assisting this group. The reaction of the Blue Leader to these violent incidents has been generally viewed as heavy-handed, and even inappropriate. The Blue Leader has branded the separatists as bandits, and has sent the worst behaving police from the north (BlueLand) to handle all protesters in the YellowLand. There are many accounts of police brutality and civilian deaths. In December 2004, a natural disaster hit and ravaged portions of YellowLand. The massive arrival of relief workers lead to an interruption of hostilities, but these resumed in mid-2005, and Blue Leader declared martial law over YellowLand in the summer of 2005.

The violent incidents in the country were classified based on the size and intensity of the incident. The incidents were then aggregated and plotted against time to obtain a longitudinal plot of incidents (Figure 9). The data was then longitudinally separated into two 'independent sets' of training and rest data. The training set consisted of Jan-June 2004, while the test set covered the period July 2004 till Dec 2004. We curtailed the test data to end before the above mentioned natural disaster.

Setting Up The Test-Bed and Tuning It With The Training Dataset: Training data and evidence were used to calibrate three types of agents in PMFserv:

- Blue Leader (structure of his GSP trees are in Fig 3) - data indicates harsh, cruel, task, corrupt, wealthy, successful. Sends worst behaving cops down to YellowLand, never discourages brutality.
- Moderate villagers -Lack of cultural freedom, schools, etc. Want own land and autonomy.
- Radical villagers - Wahhabi and college-trained, unemployed, running religious schools in family homes.

Earlier Figure 5A shows such a male and his GSP trees.

In order to adequately test these PMFserv agents' ability to interact at the population level, Lockheed Martin/ATL connected a version of our Red/Yellow/Blue scenario to a cellular automata that is known as the Civil Violence model (Epstein et al., 2001), though they replaced Leader Legitimacy with PMFserv agents' view of membership. We have since improved and altered that setup as described here. The Civil Violence model involves two categories of actors, namely villagers (or simply agents) and cops. 'Agents' are members of the general population of YellowLand and may be actively rebellious or not, depending on their grievances. 'Cops' are the forces of the BlueLand authority, who seek out and arrest actively rebellious agents. The main purpose of introducing the Civil Violence model is to provide a social network for the cognitively detailed PMFserv villagers to interact with. The social network consists of one layer of the normal arena or neighborhoods as well as a second layer of secret meeting places, simply represented as a school. Civil Violence agents can exist in more than one layer (namely in the normal as well as school layers), however, the PMFserv agents that show up in the school layer are only the young Wahhabi- and college-trained males.

The training data set also was used to fit the between-the-models parameters, especially between the PMFServ and CV model bridge and to tune up the Civil Violence villagers. Specifically, three types of cellular automata villagers were added:

- Neutral Villagers (these are modeled as simple agent automata in the CV model) -- 1,360 of them exist. The simple villagers are uniformly distributed in terms of risk aversion, but derive their grievance from witnessing cop activities in their neighborhood, from polling neighbors for opinions, and from hearing about hardships and news from PMFserv agents they may be in contact within their own neighborhood or school.
- Moderate Villagers – there are 80 of these in Civil Violence who are controlled by 80 PMFserv agents. They influence neutrals via small world theory in different neighborhoods of the Civil Violence cellular automata.

- Radical Villagers – there are 80 of these in Civil Violence who are controlled by 80 PMFserv agents. They influence neutrals via small world theory in different neighborhoods of the civil violence cellular automata and in the school layer.

The bridge between PMFserv and Civil Violence includes Blue Leader and 160 villagers, and works as follows. Blue Leader examines the state of the world and makes action decisions to assist or suppress Red (insurgents) or Yellow (Muslim community members) (e.g., pay for Buddhist schools, add more cops, reduce cop brutality, etc.). The 160 PMFserv agents then assess their view of the world, react to how cops handle protester events, how their GSPs are being satisfied or not by leader actions, and to their emotional construals. The grievance level and group membership decisions by 160 archetypical villagers in PMFServ are passed via an XML bridge to 160 agents they control in the cellular automata based population model. These agents influence the neutrals of the population who spread news and form their own view of the situation. The number of Civil Violence villagers in each of the five states of the Grievance Scale (neutral through Fight Back) are added up and this information is passed back to PMFserv to help determine its starting level of grievance for the next cycle of reactions to Blue Leader actions. For the purposes of this paper, the Red Group has no active agents, but is marked up as rebels that afford activations as mentioned in Section 3.

Running the Simulation -- The correspondence test is whether the overall parameterization for the GSP tree-guided PMFserv agents in the bridge with the Civil Violence population will faithfully mimic the test data set. That is, by tuning the GSP trees of 1 leader and 160 villagers, and by connecting all that to the Civil Violence mode of spreading news and grievances, do we wind up with a simulation that seems to correspond to what happened in the real world test dataset? Specifically, we are interested in testing the null hypothesis that there is no statistically significant correlation between real decisions and the simulated decisions. That is to say that real incidents and simulated base case are mutually independent. This assessment of modeling results will be carried out with respect to the single, actual history, and thereby ignoring any potential counterfactuals. Validation involving counterfactuals, in addition to real history, will require a lengthy assessment, which is beyond the scope of this paper, but can be found in Bharathy (2006).

The simulation starts on the left side of Figures 7A-D for Lsim and 9 for Civil Violence. When the simulation is run, one observes Blue Leader trying some assistance measures initially (usually offering to set up

Buddhist school and institutions) but maintaining a high police presence, and turning increasingly suppressive as the run proceeds -- Suppressing by Increasing Militarization and by Increasing Violence Unleashed. The end state is reflected in Figures 8B and 8C for each view, respectively. We can also examine what happened as the run proceeded. Figure 10A shows the average PMFserv villager perceptions of the Blue Leader actions in terms of the Dangerous Ideas model's terms – Vulnerability, Injustice, Distrust. Initially, Moderate Villagers respond positively to needed assistances given by the Blue Leader (negative VID and grievance is positive support). However, once they are suppressed violently and lose faith in the government, they tend to disagree with even positive government decisions. Radical Villagers start out disagreeing with Blue Leader and shift to 'fight back', an action that might continue for a long time before they realize the helplessness of the situation and abandon membership in the moderate side, and join the opposition.

Figure 10B shows the output of the Civil Violence model being sent back to the PMFserv villagers. Specifically, it shows what percent of the population has been shifted from Neutral Grievance to higher states (recall the scale of earlier Section 3): GS0 (neutral) through GS4 (fight back). From the first graph, it can be seen that at the start, most villagers are neutral and occupy GS0 while a small percent start in GS1. Many of them rapidly shift to GS1 (disagree), then abandon that and shift to GS2 and higher states. The occupancy in lower grievance states fall with time, while that in higher grievance states climb. From about week 50 onwards, there is a fairly stable, though regularly punctuated equilibrium in which the highest occupied states are GS3 and GS4. This is an indication of progressive escalation of violence in the society.

In order to compare this simulated grievance to that of the real world, we need some reliable measures of the population's grievance during actual events. Unfortunately, there are no survey or attitude results available. In the real world (test) dataset, the incident data was available, however, with a record of fatalities and injuries. There are a number of schemes for weighting those (e.g., depression and morale loss, lost income, utility metrics, others), however, here we take the simple approach of just computing a weighted incident severity. We computed incident severity scores using weighted average of fatalities and injuries, where injuries are simply counted, but the weight on fatalities is 100.
$$IncidentSeverity = w_f \times fatalities + w_i \times injuries \quad [6]$$

The result serves to indicate how severe these incidents were. While severity is only an indirect measure of how the population might have felt, it is a measure that can be tested for correlation to the rise and fall of grievance expression due to leader actions in our simulated world.

To conduct the comparison, we apply the non-parametric Kendall's Tau measure of correlation (Hollander and Wolfe, 1999). This statistic estimates the excess of concordant over discordant pairs of data, adjusted for tied pairs. With a two sided test, considering the possibility of concordance or discordance (akin to positive or negative correlation), we can conclude that there is a statistically significant lack of independence between base case simulation and observed grievances rankings at a confidence interval of 88%. Since there is a probabilistic outcome determining if a simulated leader's action choice will result in injury and fatality incidents (and how the news of these events are propagated through the cellular automata is probabilistic as well), we repeated the simulation runs thirty times and the confidence interval mentioned above is the mean across those 30 correlations. In sum, the null hypothesis is rejected and real (test interval) incident data and simulation results are related.

Figure 10A – PMFserve Villagers View the Blue

Figure 10B – Percent of Civil Violence Villagers Shifting

2. Turing Test

The Turing test looks beyond population statistics, and examines what transpires inside the heads of the various types of agents in the simulated world. The Turing Test of Sect 4.1.2 showed the PMFserv emulation of Crusades leaders to be reasonably credible. In this case study, we have more detailed data and apply a more rigorous test. Specifically, in the test dataset, the real world leader made 52 decisions affecting the population and that we sorted into positive, neutral, and negative actions. In the simulated world, Blue leader made 56 action decisions in this same interval. At this level of classification (positive, neutral, negative), we were able to calculate a mutual information or mutual entropy (M) statistic between the real and simulated base cases. M ranges from 0 to 1.0, with the latter indicating no correlation between two event sets X and Y.

$M(X:Y) = H(X) - H(X|Y)$ where X and Y are the simulation and historic sources, respectively, and H(.) is the entropy function, defined by: $H(X) = - \sum p(x)_i \log p(x)_i$. Applying this metric, the mutual entropy values were found to be less than 0.05 (at least an order of magnitude smaller than the mutual entropy of 1.0), indicating reasonable degree of correlation between real and simulated data. With an M metric, one cannot make statements about the confidence interval of the correlation, however, like the Crusades leaders, the Blue Leader in the current scenario

seems equally faithful to his real world counterpart. This gives us reason to suspect that the Hermann- and GLOBE-based GSP tree structure may work equally well across time periods, locations, and cultures.

V. LESSONS LEARNED AND NEXT STEPS

In concluding, it is useful to revisit the four aims of the introduction, and to see what has been learned in each of them and to point out some items seen as priorities for further development.

A. Socio – Cultural Game Generator (Aim 1)

Aim 1 was to create a role-playing game generator where one could rapidly set up and play out numerous conflict scenarios from around the world. Conflicts arise when groups vie over the control and allocation of resources (land, economy, markets, militias, media outlets, followers, etc.). Socio-cultural aspects concern any perceived injustices that have arisen historically with respect to these allocations, where perception is a matter of the value systems, norms/standards, and emotional utility of the perceivers. Lsim implements such a conflict scenario generator and this article showed two examples of its usage: Christians versus Muslims in the Crusades, and Buddhists versus Muslims in modern day South East Asia. The game generator was shown to reduce conflicts to the bare essentials that allow one to explore the intertwined issues affecting welfare (economy, in-group standards, health services), security (freedoms/liberties, military), and political support for leaders (popularity of positions).

In zero sum games, what one spends on actions affecting one area of welfare, security, or populace effects what one has to allocate to other areas. Borrowing from diplomatic video games the idea here is to make the game immersive and engaging, and to date hundreds of players have participated in multi-hour sessions that they were unwilling to terminate. All this game-play also gave us a rich source of data to help guide the construction of agents who can serve as synthetic opponents, allies, followers, and the like. Also, we have learned that our game state representations are intuitive and that domain experts can readily use them to express conflict scenarios that are hard to verbalize. As with anything done in software, there are always next levels of sophistication and detail that one can add, and we identified many new features we would like to add such as, to mention a few examples, (1) scale up of all features shown here for the larger game generator we call Athena's Prism; (2) resources and assets (e.g., economy and black markets) that are supported by institutions that grow more self-sustaining and resilient, the larger

they are; and (3) logging services and explanation functions that help users to generate reports on model outcomes, agent decision choices, and effects. These are some of the laundry list of next steps for the game generator.

B. Using “First Principles” and Synthesizing Social Science (Aim 2)

Modeling leaders and followers is a complex enterprise and one would like to use only first principles of social science, yet the field has not matured sufficiently. Still, that is no excuse for modelers to “make up” their own rules and algorithm for how groups behave, nor is it justification to just create entertaining agents. The alternative we explored here is to try and adopt best-of-breed and well-respected social science models for leadership, group dynamics, and the hearts and minds of the populace (AIM2). These models are implemented atop a unified architecture of cognition, call PMFServ that manages six modules of an agent’s mind: memory, perception, physiology/stress/coping level, value system, and emotional construal, relationships and models of other, and (stress and emotion-constrained) decision making processes. PMFServ exposes many parameters in each of these modules and permits analysts/developers to visually “program” best-of-breed social science models that govern how the modules work, and in turn, how that agent tends to behave. This framework supported the ready implementation of leader models from Hermann (style), Hofstede and Globe (cultural factors), and Heuer (biases) atop pre-existing models in the PMFServ modules. These synthetic leaders passed the Turing and Correspondence tests in both conflict scenarios—The Crusaders and SE Asia – where each leader attempted to maximize his respective economic welfare, security, and populace resources in accord with his GSP trees of goals, standards, and preferences in the game scenarios. It was no surprise that leaders’ biased models of other leaders often proved to be self-fulfilling prophecies. It was a surprise; however, that once leaders satiated their main preferences they chose “cooperation solutions” if the opponent leader chose to specialize in non-threatening resource allocation choices.

Likewise, the PMFServ modules allowed group followers to be readily modeled via their personal motivations (Maslow-style), group member factors (injustices, vulnerabilities, etc.), and loyalty decisions (follow happily, helplessly, vocally, separate, etc.). Again the followers’ behavior passed Turing and Correspondence tests of Muslim moderates and radicals as the outgroup leader’s policies shifted: a real world case study was used. Our population model involved a cellular automata with 1,360 agents influenced in their neighborhoods and schools by 160 PMFServ agents. Inside the PMFServ agents, one can readily observe and track their GSP tree implementation of Maslow, Hofstede, and GLOBE factors, and preference functions. One can follow how they update the Eidelson

model factors of group and leader achievement. One gains confidence that these agents are realistic, particularly when one can calibrate them with validated instruments such as Hermann's profiling method or Eidelson's IGBI instrument, just as is done for real world human participants: e.g., see Moaz & Eidelson (2006).

C. Improving the Science (Aim 3)

Our contention is that the state of social science today can be advanced by computational rigor such as we have sought to impose. That is, when well-respected models are implemented computationally in socio-cultural games, they are more fully tested, their limitations exposed, and the needs for improvement and refinement are delineated. As one example of this, in order to use the Hermann leader profile instrument we first had to derive a mathematics to quantify each factor relative to game actions and state (e.g. grow vs. protect, task vs. relations, in-group bias, etc.). In PMFserv terms this means mapping the Hermann factor sets to a leader's GSP trees and activation mechanics so that it becomes a subjective (emotional) expected utility implementation. In trying to build leaders that function in the game, however, the Hermann factors proved to be not quite sufficient, and we had to quantify and add some factors of Hofstede and from the UN's GLOBE study. The aim of our research is realism not mathematical elegance, however, and this means we also added Heuer-suggested biases and heuristics such as mirroring and confirmation, among others. The end result worked well enough on Turing and Correspondence tests for us to be able to make the following claim about leader modeling: a synthesis of respected models is important if the science and practice are to advance. Also, a virtual laboratory is essential to foster the syntheses needed in social sciences.

This paper presented a second example, that of group member modeling, where we explored unification of a motivation model (Maslow) with a group profiling instrument (Eidelson's Dangerous Ideas or DI model) and the Hirshman model of loyalty, voice, and exit. Mathematically these were once again cast into our subjective (emotional) expected utility formalisms and implemented within PMVserv's six modules and the other best-of-breed models already existing inside PMFserv and the game. The 4th co-author of this paper is in fact the developer of one of these three group modeling approaches and we report his comments on whether the experience of participating in socio-cultural game improved the science and/or practices:

"Implementing my dangerous ideas framework (Eidelson & Eidelson, 2003) in the PMFserv framework produced a series of challenges and opportunities, and required that I substantially advance my own thinking in

order to effectively address the many operationalization issues that arose. In particular, the process of making underlying assumptions transparent and making key relationships among variables quantifiable in response to questions from my modeling collaborators led me to critically evaluate and refine both abstract and specific components of the framework. In short, this enterprise has definitely helped to move my own theoretical analysis forward.”

- Roy Eidelson, 2006

This seems to support our claim that science can be advanced by studying its strengths and weaknesses in socio-cultural games. In fact, we contend this is a vital pathway to improving the science of this field in general.

D. Analyzing Alternatives (Aim 4)

The entire point of insisting on well-respected models inside and on validation efforts for the synthetic agents is so one can have trust that experiments on these agents will yield insights about the alternative policies that influence them. As mentioned above, a major objective of this research is to utilize experiments on synthetic agents to identify those policy instruments that will most influence the real-world agents they represent. Page limits prevent us from presenting the numerous sensitivity experiments we have run and these will appear in future papers.

ACKNOWLEDGEMENT: This research was partially supported by the US Government (game engine, leader models), DARPA (follower models, 2nd case study), AFOSR (correspondence tests, sensitivity studies), and by the Beck Scholarship Fund. Lockheed Martin/ATL prepared an early version of the cellular automata framework we updated and used here. Also, we thank the US Government for their guidance and encouragement, though no one except the authors is responsible for any statements or errors in this manuscript.

REFERENCES

- [1] Green and Armstrong (2003), J.S. Armstrong. (2002). Assessing game theory, role playing and unaided judgment. *International Journal Forecasting*, 18, 345-352

- [2] R. J. Heuer, Jr., *Psychology of Intelligence Analysis*. Washington, DC: Center for the Study of Intelligence, Central Intelligence Agency, 1999.

- [3] M. van Lent, R. McAlinden, P. Brobst, et al., “Enhancing the behavioral fidelity of synthetic entities with human behavior models,” presented at the 13th Conference on Behavior Representation in Modeling and Simulation (BRIMS), Simulation Interoperability Standards Organization (SISO), Arlington, Virginia, 2004, pp 12-15.
- [4] B.G. Silverman, K. O’Brien, J. Cornwell (2006b, April). Human Behavior Models for Agents in Simulators and Games: Part II – Gamebot Engineering with PMFserv. *Presence*. 15: 2.
- [5] D. McDonald, R. Lazarus, A. Leung, et al., “Interoperable Human Behavior Models for Simulations,” presented at the 15th Conference On Behavioral Representation in Modeling & Simulation (BRIMS), SISO, May. 2006
- [6] B.G.Silverman, R. Rees, J. Toth, et al. “Athena’s Prism: A Diplomatic Strategy Role Playing Game for Generating Ideas and Exploring Alternatives,” presented at Conference on Intelligence Analysis, MacLean, VA: Mitre, 2005.
- [7] B.G Silverman, M. Johns, R. Weaver, K. O’Brien, R. Silverman (2002a). Human behavior models for game-theoretic agents. *Cognitive Science Quarterly*, 2 (3/4), pp. 273-301.
- [8] B. G. Silverman, M. Johns, K. O’Brien, R. Weaver, J. Cornwell, “Constructing virtual asymmetric opponents from data and models in the literature: Case of crowd rioting,” *Proceedings of the 11th Conference on Computer Generated Forces and Behavioral Representation*, 2002b, pp. 97-106, Orlando, Florida, 97-106.
- [9] B.G. Silverman, (2005). Human performance simulation. In J.Ness, D.Ritzer, & V.Tepe (Eds.), *The Science and Simulation of Human Performance (Chapter 9)*. New York: Elsevier.
- [10] B.G. Silverman, M. Johns, J. Cornwell, K. O’Brien, “Human Behavior Models for Agents in Simulators and Games: Part I – Enabling Science with PMFserv.” *Presence*, Vol. 15, No. 2, Pages 139-162, April 2006,

- [11] M. Woolridge, P.E. Dunne, "The computational complexity of qualitative coalitional games." *Artificial Intelligence*.
- [12] G.I. Simari, S. Parsons, "On approximating the best decision for an autonomous agent." Paper presented at the Third International Joint Conference on Autonomous Agents & Multi Agent Systems (AAMAS), W16: Workshop on Game Theory and Decision Theory. New York City, July 19-23, 2004.
- [13] R. W. Pew, A. S. Mavor, *Modeling human and organizational behavior: Application to military simulation*. Washington, DC: National Academy Press.
- [14] M. M. Chemers, *An Integrative Theory of Leadership*. Mahwah, NJ: Lawrence Erlbaum Associates.
- [15] M. G. Hermann. (1999). Assessing leadership style: A trait analysis. *Social Science Automation, Inc*.
- [16] R.J. House, P.J. Hanges, M. Javidan, et al., *Culture, Leadership, and Organizations: The GLOBE Study of 62 Societies*, Thousand Oaks, CA: Sage Publications
- [17] Bharathy (2006) *A knowledge engineering based systems methodology for integrating of social science frameworks for modeling agents with cognition, personality & culture*. (Doctoral dissertation, University of Pennsylvania, 2006)
- [18] A.R. Damasio, *Descartes' error: Emotion, reason, and the human brain*. New York: Avon.
- [19] A. Ortony, G. L. Clore, A. Collins, *The Cognitive Structure of Emotions*. Cambridge, MA: Cambridge University Press, 1988.
- [20] R. Lazarus, *Emotion and Adaptation*. Oxford: Oxford University Press, 1991

- [21] M. Johns, *Deception and Trust in Complex Semi-Competitive Environments* (Doctoral dissertation, University of Pennsylvania, 2006)
- [22] G. Hofstede. (2003). *Culture's Consequences, Comparing Values, Behaviors, Institutions, and Organizations Across Nations*. Newbury Park, CA: Sage Publications. [Online]. Available: <http://www.geert-hofstede.com/>
- [23] M. Sageman, *Understanding Terror Networks*. Philadelphia, PA.: University of Pennsylvania, 2005
- [24] R. J. Eidelson, J.I. Eidelson. (2003). Dangerous ideas: Five beliefs that propel groups toward conflict. *American Psychologist*. 58, pp. 182-192.
- [25] J. Cornwell, K. O'Brien, B.G. Silverman, J. Toth, "Affordance Theory for Improving the Rapid Generation, Composability, and Reusability of Synthetic Agents and Objects," presented at the 12th Conference Behavior Representation in Modeling and Simulation (BRIMS, formerly CGF), SISO, 2003.
- [26] B.G. Silverman and G. Bharathy. (2005, May). "Modeling the Personality & Cognition of Leaders." Presented at the 14th Conference on Behavioral Representations In Modeling and Simulation, SISO, Available: <http://www.sisostds.org/>
- [27] T. Hobbes, *Leviathan*. New York: E. P. Dutton, 1947 [1651]
- [28] R. Axelrod, D.S. Bennett. (1993). A landscape theory of aggregation. *British Journal of Political Science* Vol. 23, pp. 211-233. [Reprinted in Axelrod, R. (1997). *The Complexity of Cooperation*. Princeton, NJ: Princeton.]
- [29] A. J. Field, *Altruistically Inclined? The behavioral sciences, evolutionary theory, and the origin of reciprocity*. Ann Arbor, MI: T. Kuran (Ed.), *Economics, Cognition, and Society*, University of Michigan Press.

- [30] Mamoun Fandy. (1994). Tribe versus Islam: The Post-Colonial Arab State and the Democratic Imperative. Middle East Policy Council. 3, pp 40-51.
- [31] J. Epstein, J.D. Steinbruner, M.T. Parker, "Modeling Civil Violence: An Agent-Based Computational Approach," Proceedings of the National Academy of Sciences. Washington DC: Brookings (CSED WP#20), 2001
- [32] M. Hollander, D.A. Wolfe, *Non-parametric statistical methods*, 2nd ed., New York: Wiley.
- [33] I. Maoz, R.J. Eidelson. (2006). Psychological Bases of Extreme Policy Preferences: How the Personal Beliefs of Israeli-Jews Predict Their Support for Population Transfer in the Israeli-Palestinian Conflict. *American Behavioral Scientist*, accepted.
- [34] JE Cote, E James, C Levine (2002). *Identity Formation, Agency, and Culture*. New Jersey: Lawrence Erlbaum Associates.

Figure 1 – Overview of the Basic Leader-Follower Game That is Repeated Around the World

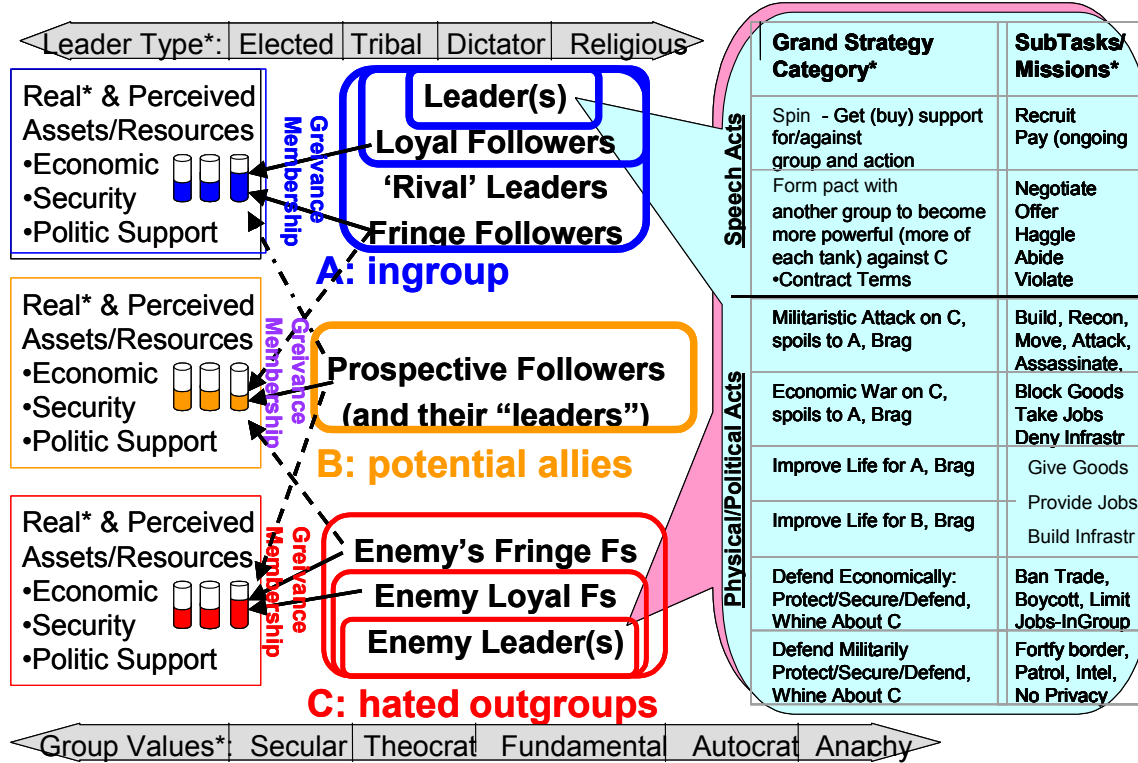








Figure 2 – LeaderSim: a World Diplomacy Role Playing Game for Humans or Agents

LeaderSim AI Prototype
Game Actions Intent Options Help





RedLand

Economy 
 People 
 Armed Forces 

BlueLand

People 
 Armed Forces 
 Economy 

YellowLand

People 
 Economy 
 People 
 Armed Forces 


Action Target Resource


Action Reference



Hostile Actions


Name	Funded by	Targets	Defended by
Conventional Attack	Armed Forces	Armed Forces, People	Armed Forces
Economic Attack	Economy	Economy, People	Economy
Asymmetric Attack	People	Armed Forces, Economy, People	Armed Forces, People

Economic Attack Result

 Blue instigated a Economic Attack action with Power of 1




 Red defended against a Economic Attack with Power of 1

Attacker	Die Rolls	Defender
 Power: 1	5 2	 Power: 0




 Blue wins by 1!

LeaderSim AI Prototype
Game Actions Intent Options Help





RedLand

Economy 
 People 
 Armed Forces 

BlueLand

Armed Forces 
 People 
 Economy 

YellowLand

People 
 Economy 
 Armed Forces 
 People 

Non-Hostile Actions

- Manage Domestic Affairs: Renew any five resource units
- Move: Transfer units of one resource into another
- Give Foreign Aid: Spend resources to allow another player to renew resources

Figure 3 – GSP Tree Structure, Weights and Emotional Activations for Country Leader

PMFServ Integrated Development Environment
 File Library Scenario View Window Help

New Scenario

- Agents
 - Guy
 - Montferrat
 - Philip
 - Richard
 - Richard's GSP Tr
 - Stress
 - Physiology
 - Relationships
 - State
 - Action Bindings
 - Perceptual Types
 - Decision
 - Emotions
 - Perception
 - Saladin
 - Saladin's GSP Tr
 - Stress
 - Physiology
 - Relationships
 - State
 - Action Bindings
 - Perceptual Types
 - Decision
 - Emotions
 - Perception
- Objects
 - Simulation

Detail Panel - Saladin's Emotions

Goal Emotions: Fortunes of Self

Distress Joy

Fear Hope

Preference Emotions

Disliking Liking

Goal Emotions: Fortunes of Others

Resentment Pity

Gloating Happiness

Standard Emotions: Attributions

Shame Pride

Reproach Admiration

Anger Gratitude

Remorse Gratification

ActorId	Action	Result	Target Id
Saladin	Perceive	Perceive	Saladin
Richard	Perceive	Perceive	Saladin
Philip	Perceive	Perceive	Saladin
Saladin	Perceive	Perceive	Saladin
Saladin	Perceive	Perceive	Richard
Saladin	Perceive	Perceive	Philip
Saladin	Perceive	Perceive	Montferrat
Saladin	Perceive	Perceive	Guy

Goals

- 0.40 — Grow
- 0.60 — Protect
- Standards**
 - 0.33 — Domestic_Policy
 - 0.33 — Exercise_of_Power
 - 0.40 — Be_Controlling
 - 0.60 — Be_Open
 - 0.33 — Scope_of_Doing_Good
 - 0.70 — Bring_About_Greater_Good
 - 0.30 — Look_After_Narrower_Interests
 - 0.34 — Task_Relationship_Balance
 - 0.41 — Be_Relationship_Focussed
 - 0.59 — Be_Task_Focussed
 - 0.33 — Foreign_Policy
 - 0.60 — Diplomacy
 - 0.33 — Provide_Help
 - 0.10 — Help_Enemies_too
 - 0.65 — Help_Friends
 - 0.25 — Help_Neutrals
 - 0.67 — Resolve_Issues_by_Negotiation
 - 0.40 — Military
 - 0.15 — Use_Assymmetric_Attacks
 - 0.80 — Use_Conventional_Attacks
 - 0.66 — Use_ONLY_Conventional
 - 0.34 — Use_also_Conventional
 - 0.34 — Treatment_of_Out_Groups
 - 0.50 — Out_Groups_are_Legitimate_Targets
 - 0.50 — Enemy_is_Out_Group
 - 0.10 — Friend_is_Out_Group
 - 0.40 — Neutral_is_Out_Group
 - 0.50 — Treat_with_Fairness_n_Justice
 - 0.67 — Life_Res_r_Sensitive
 - 0.33 — None_r_Sensitive
 - Preferences
 - 0.30 — Contested_Territory_CT
 - 0.20 — Ally's resources in CT
 - 0.25 — Enemy's resources in CT
 - 0.15 — Neutral's resources in CT
 - 0.40 — Self's resources in CT
 - 0.20 — Enemy_Territory_ET
 - 0.10 — Friendly_Territory_FT
 - 0.05 — Neutral_Territory_NT
 - 0.35 — Own_Territory_OT

Figure 4 – Screens from Lsim and PMFserv Depicting a Buddhist-Muslim Conflict in SE Asia

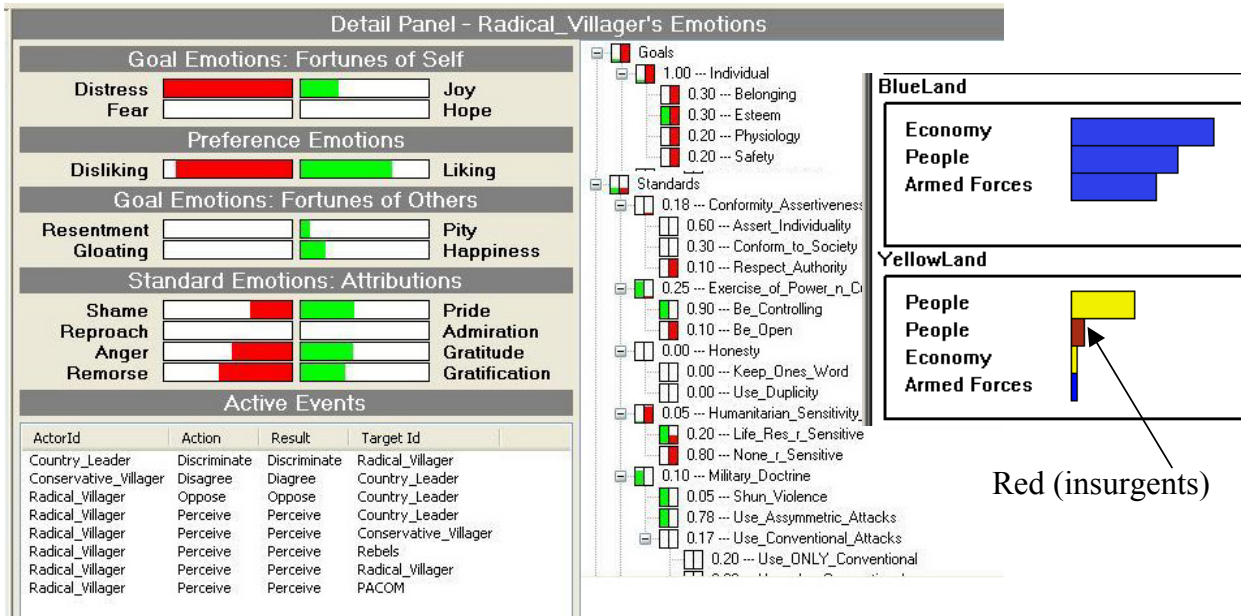


Figure 5 – Example of a Villager’s Range of DI Perception Possibilities When Viewing a Given Leader




Figure 6

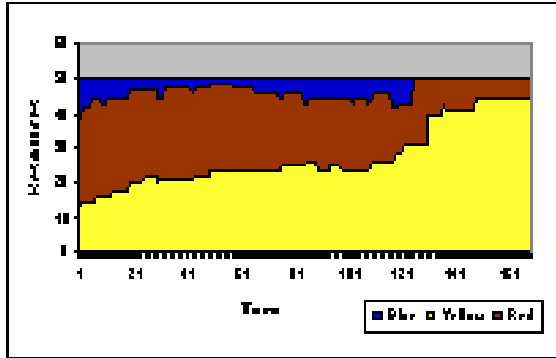


Figure 6A: Resources vs. Turns for Outcome I




Figure 6B: Resources vs. Turns for Outcome II

Figure 7

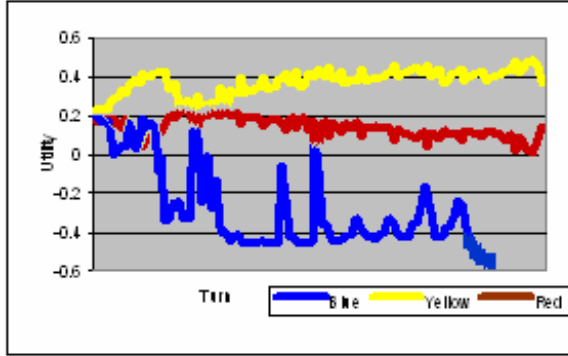


Figure 7A: Utility vs. Time

Yellow's utility is highest since he is able to grow. Blue is driven to the brink of extinction early on, and suffers a big dip in utility as a result. Red's utility of the game shrinks only marginally. This may be since his GSP is task-focused (war itself) and insensitive to life, plus he is able to withdraw with a symbolic win (People remain).




Figure 7B: Impression Error vs. Time

By mirroring, Red is consistently able to assess the compatible GSPs of Blue, but Red suffers from significant error when it comes to assessing the goals and preferences of Yellow.




Figure 7C: Power vs. Time

The above figure shows the rise of power for Yellow and the fall of Blue and Red. These parameters activate "distress vs. joy" on the leaders' Goal Tree, as well as the emotions about the fortunes of others (e.g., gloating, resentment, etc.).


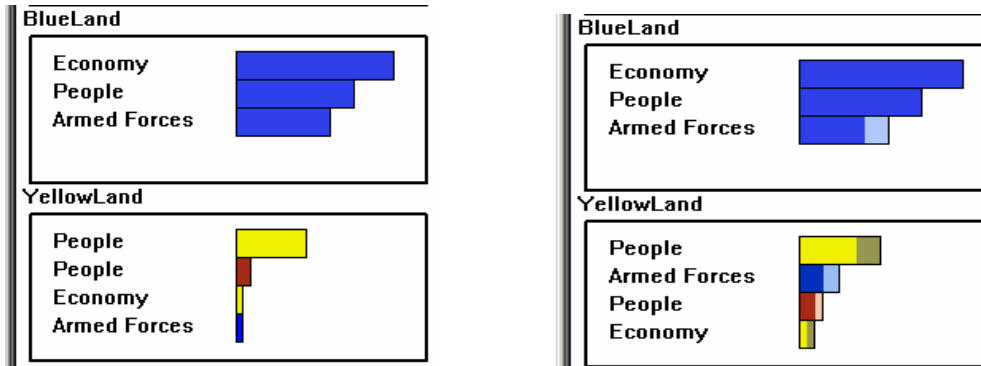


Figure 7D: Preference Failure vs. Time

Failure (1-preferences) measures how unsatisfied each of the leaders is over time. Yellow-L's low and declining preference is understandable in a world of growth. This measure correlates with the "like and dislike" emotions.

Figure 8 – Start and End States During the CorrespondenceTest: LeaderSim Summary View



A – Blue Land at Start of Test with Minimal Forces in Yellow Land, Minimal Insurgency

B – By End of Test, Blue Leader has Moved Forces to Yellow Land, Insurgency Grows

8C - Civil Violence View of Population Membership Starting and Ending States

Grievance State	C – Starting State (Avg of Weeks 1 & 2)	D – End State (Avg of Weeks 103, 104)
GrievanceState0 - Neutral	30%	6%
GrievanceState1 - Disagree	55%	1%
GrievanceState2 - Join Opposition	15%	37%
GrievanceState3 - Nonviolent	0%	39%
GrievanceState4 - Fight-Rebel	0%	17%
TOTAL	100%	100%

Lsim Legend: Blue = Leader and Cops/Armed Forces, Yellow = Moderates and Radicals, Red = Rebels

Figure 9 –Violence Events on Grievance Scale




Figure 10

Figure 10A – PMFserv Villagers View the Blue Leader’s Actions And Communicate That to Their Civil Violence Counterparts




Figure 10B – Percent of Civil Violence Villagers Shifting from Neutral To Higher States of Grievance




Table 1 – The Seven Traits of the Hermann Leadership Style Profile

1. Belief that one can control what happens	Combination of the two attributes (1) and (2) determines whether the leader will challenge or respect the constraints.
2. Need for power/influence	
3. Concept complexity	Combination of the two attributes (3) and (4) determines how open a leader will be to information.
4. Self-confidence	
5. Task Vs Relationship Focus	a continuum between two poles: <ul style="list-style-type: none"> ○ Moving the group toward completion of a task, and ○ Maintaining group spirit and morale (building relationships).
6. An individual's general distrust or suspiciousness of others	The leader's outlook about the world and problems largely determines the confrontational attitude of the country, likelihood of taking initiatives and engaging in sanctions. The extent of their in-group bias and general distrust of others is driven by: <ul style="list-style-type: none"> ○ perceived threats or problems in the world, or ○ perceived opportunities to form cooperative relationships.
7. The intensity with which a person holds an in-group bias.	

Table 2 – The Grievance Scale Embedded in Follower Agents

← (in support of the Country Leader) Villager Decision (against the Country Leader) →								
Sacrifice, Go on Attacks	Support, Vote for	Join Authority	Agree	Neutral	Disagree, Vote against	Join, Opposition	Oppose, Non- Violent	Fight Rebel,
-4.0	-3.0	-2.0	-1.0	0.0	+1.0	+2.0	+3.0	+4.0
GS0					GS1	GS2	GS3	GS4