



August 2006

Gaming and Simulating EthnoPolitical Conflicts

Barry G. Silverman

University of Pennsylvania, basil@seas.upenn.edu

Gnana Bharathy

University of Pennsylvania

Benjamin Nye

University of Pennsylvania

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

Recommended Citation

Barry G. Silverman, Gnana Bharathy, and Benjamin Nye, "Gaming and Simulating EthnoPolitical Conflicts", . August 2006.

Postprint version. Presented at *Descartes Conference on Mathematical Models in Counterterrorism*, September 2006, 19 pages.

This paper is posted at ScholarlyCommons. http://repository.upenn.edu/ese_papers/288

For more information, please contact libraryrepository@pobox.upenn.edu.

Gaming and Simulating EthnoPolitical Conflicts

Abstract

This paper begins by describing a universally recurring socio-cultural 'game' of inter-group competition for control of resources. It next describes efforts to author software agents able to play the game as real humans would – which suggests the ability to study alternative ways to influence them, observe PMESII effects, and potentially understand how best to alter the outcomes of potential conflict situations. These agents are unscripted, but use their decision making to react to events as they unfold and to plan out responses. For each agent, a software called PMFserv operates its perception and runs its physiology and personality/value system to determine fatigue and hunger, injuries and related stressors, grievances, tension buildup, impact of rumors and speech acts, emotions, and various collective and individual action decisions. The paper wraps up with a correspondence test from a SE Asian ethnic conflict, the results of which indicate significant correlation between real and agent-based outcomes.

Keywords

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

Comments

Postprint version. Presented at *Descartes Conference on Mathematical Models in Counterterrorism*, September 2006, 19 pages.

Gaming and Simulating EthnoPolitical Conflicts

Barry G. Silverman, Gnana Bharathy, Benjamin Nye

Electrical and Systems Engineering Dept.

University of Pennsylvania, Philadelphia, PA 19104-6315

basil@seas.upenn.edu

August 2006

Abstract

This paper begins by describing a universally recurring socio-cultural ‘game’ of inter-group competition for control of resources. It next describes efforts to author software agents able to play the game as real humans would – which suggests the ability to study alternative ways to influence them, observe PMESII effects, and potentially understand how best to alter the outcomes of potential conflict situations. These agents are unscripted, but use their decision making to react to events as they unfold and to plan out responses. For each agent, a software called PMFserv operates its perception and runs its physiology and personality/value system to determine fatigue and hunger, injuries and related stressors, grievances, tension buildup, impact of rumors and speech acts, emotions, and various collective and individual action decisions. The paper wraps up with a correspondence test from a SE Asian ethnic conflict, the results of which indicate significant correlation between real and agent-based outcomes.

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

1) 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 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.

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 generation and testing of new hypotheses is a third aim, and another potential benefit of simulations.

1.1) 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 fairly 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 mock up such conflict scenarios and analyze what outcomes arise from different courses of action/policies. We describe this game intuitively here and more formally in Appendix I.

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. Before delving into our model, we invite the reader to take a look at a stylized resource control game through game theoretic framework, as set up in the Appendix 1. The inspiration for this stylized game comes from Wood's (2003) civil war settlements paper. This game theoretic exercise not only gives you a 5000ft view, but also illustrates why and how deep models such as ours will be useful. While a number of assumptions made by the game theoretic frameworks are defensible (well-ordered preferences, transitivity), others are meant for mathematical elegance. Without assumptions doing most of the "heavy lifting", it is impossible develop mathematically tractable models (De Marchi, 2005).

Many of these stylized game models are unable to encode domain information, particularly the depth of the social system. For example, human value systems are almost always assumed, hidden, or at the best, shrunk for the purpose of mathematical elegance. Yet, human behavior is vital to the conflict-cooperative game behavior.

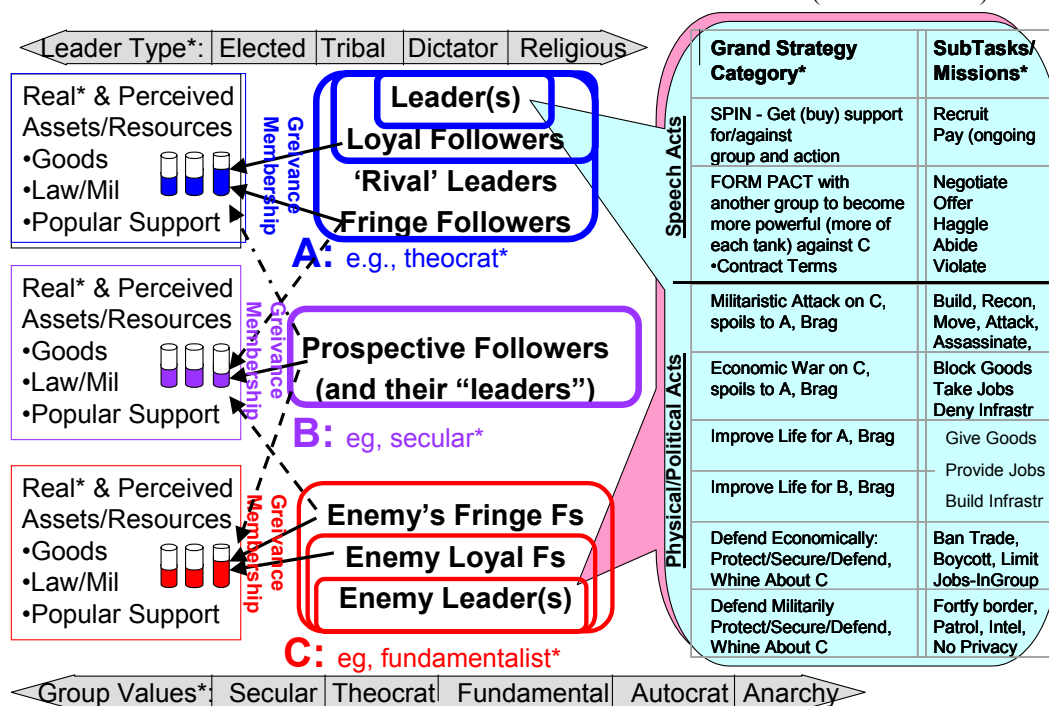
While mathematical convenience is one explanation, there is more involved. Many modeling platforms would simply not allow value systems to be made explicit, and there is no modeling process that would allow one to revisit the values. As computational power increases to accommodate more complex models, social system modelers are beginning to address this curse of simplicity.

Even though such models can not be solved mathematically, we can find solutions through validated simulation models with deep agents. If one could find clusters of parameters that pertain to a corresponding game model, we can also start talking about correspondence between game theoretic models and cognitively deep simulation models. There is room for a lot of synergy.

Now, let us return to the cognitively detailed game. 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) political goods available to the members (jobs, money, foodstuffs, training, healthcare etc.); (2) 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 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).

To start a game, there are initial alignments coded manually, though 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 outgroup 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 (spin/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.

Figure 1 – Overview of the Basic Leader-Follower Game within CultureSim (* - editable list)



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, 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© to appear at toy stores in early 2007 and 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).

In general, when humans play the game, they rapidly evolve a portfolio of strategies that they tend to pursue asynchronously and 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.

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,b, 2005, 2006a,b).

2) Agent Personality, Emotions, Culture, and Reactions

In LeaderSim, 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. 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. As described in the authors' other writings, the Standard tree implements a personality profiling tool that is mostly Hermann traits governing personal and cultural norms (Hermann, 1999), 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 and sensitivity to life (humanitarianism). 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 where many of the pitfalls associated with shortsighted AI can be sidestepped. 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 nonintuitive 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.

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. To this, Athena also adds 3 options for manage reputation (switch from none, to mirroring, to bounded rational) instead of just mirroring in Lsim.

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 write-up. As a precursor to that demonstration and to further illustrate how GSP trees represent the modified Hermann profiles, consider the right side of Figure 2. There we see the weighted GSP tree of a leader of a SE Asian nation (name withheld at request of our sponsor) who will be called BlueLeader. 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. Likewise, there are subnodes for the intersection of In Group Bias vs. Degree of Distrust. 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 BlueLeader's case, he leans heavily toward power and growth which is also consistent with his P-tree weights on his own resources. His standards reveal him to be dishonest, narrow in scope (self-interested), and task-oriented. While the figure does not expand the lower level detail of all nodes, he also is insensitive to life and outgroups (he did order the harsh treatment which lead to the slaughter of unarmed villagers).

Just to the left of the weight value on each node of the GSP trees of Figure 2 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 2 shows early in a game where BlueLeader, a Buddhist, has a lot of conflicted emotions about the outgroup of Muslim villagers in the Southern provinces and some of their demonstrations.

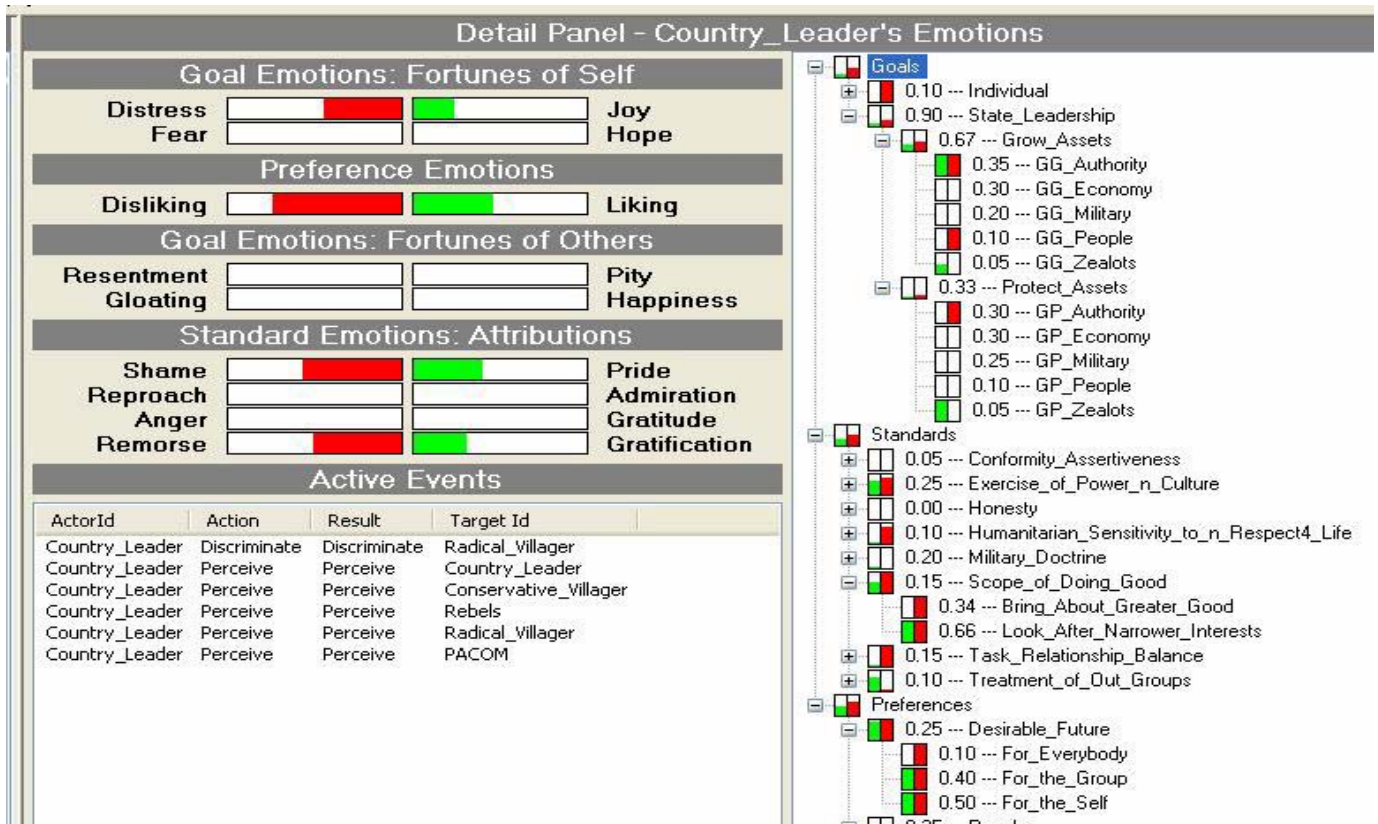


Figure 2 – GSP Tree Structure, Weights and Emotional Activations for BlueLeader

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 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 goals, standards, and preferences 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 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; etc), 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.

2.1) 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. 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. Intuitively, by factoring vulnerability into the world utility calculation, 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 vulnerability score β 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 of the current state of the world, G , when divorced from his value system and other factors:

$$UI(Gx) = \alpha x - \beta x \quad [3]$$

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. Thus α and β 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.

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

Where,

$a \in A$

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

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

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

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 will be 11 pairs of oppositely valenced activations in PMFserv's emotion model, we normalize the sum as follows so that utility varies between -1 and +1.

3) Modeling Follower Value Systems

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 (exit, voice, and loyalty). In keeping with Aim 2, each of these refinements is an implementation of a well-respected model drawn from the social sciences. Details are omitted, but may be found in Silverman, Bharathy et al (2006c).

Mathematically, the reader may recall $\phi(r_{ij})$ from earlier Equation [1]. Here we examine the case where j is a group (or leader) and the term refers to the membership, relationship, or strength of affinity of agent i to group j . An agent i can belong to multiple groups at varying strength according to:

$$\Phi(r_{iA}) = \text{Superiority}_A \times \text{GSPcongruence}_{iA} / \text{VID}_{Ai} \quad [6]$$

where Superiority and VID are from DI instruments if available, else derived by GSP trees of agent i in reacting to leader or group A.

Groups are characterized by GSP weights for the average of all members as well as by property lists defined a priori (religion, political system, etc.), and GroupPorosity factors. GSP congruence is estimated using the sum of the means square differences in the GSP nodes. $\text{GSPcongruence} = 1 - \text{Sqrt}[\text{Sum}[(w_{i1} - w_{i2})^2]]$, which is the correlation of the weights between two GSP trees. If an agent is in Group B, it will not be drawn to a Group C whose GSP archetype is substantially incongruent to its own. If an agent is in a group (or under control of a leader) whose average GSP is greatly different from its own, the agents tend to use Voice to resist the leader or attempt to Exit to another group, depending on porosity.

If agent i desires to exit from any group A to join any C, this is governed by the delta in utility of membership in each group plus a cost factor adjusted for transfer rate or demand elasticity. If the delta is positive, or larger than some loyalty factor, exit may occur. Let, this delta be:

$$\Delta\Phi_i = [U(\Phi_C) + \text{COST}_{TR} / \text{TR}_{AC}] - U(\Phi_A)$$

where,

$U(\phi)$ = utility of membership, found by invoking the PMFserv emotion model and GSP trees [7]

COST_{TR} = cost of migration, land costs, and lost opportunity costs

TR_{ij} = Transfer Rate or group porosity, a measure of ease of entry to or exit from group j for agent i .

(TR is a negative value that grows larger as porosity grows)

[8]

$$TR_{A \rightarrow C} = Salienc_{ExitA} \times Salienc_{EnterC} \times GSPcongruenc_{iC}$$

This is the transfer rate and it varies between (0,-1). Salienc is the extent to which a group permits exiting by ingroup members, and entry by outgroup members. It is the porosity permitted by the group. There is a tuple or value pair that gives both saliencForEntry and saliencForExit. The demand elasticity for exiting a group is 1/TR.

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):

Villager Decision								
←								→
Sacrifice, Go on Attacks	Support, Vote for Group A	Join Authority Group A	Agree	Neutral (undecideds in Group B)	Disagree, Vote against A	Join, Opposition Group C	Oppose, Non-Violent	Fight Rebel, Exit A
-4.0	-3.0	-2.0	-1.0	0.0	+1.0	+2.0	+3.0	+4.0
				GS0	GS1	GS2	GS3	GS4

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 exiting A and joining the insurgency. At 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 represent Voice. We only permit Exit from A and joining of C after occupying GS4 for a significant interval.

Figure 3 – Start and End States During the Correspondence Test: LeaderSim Summary View

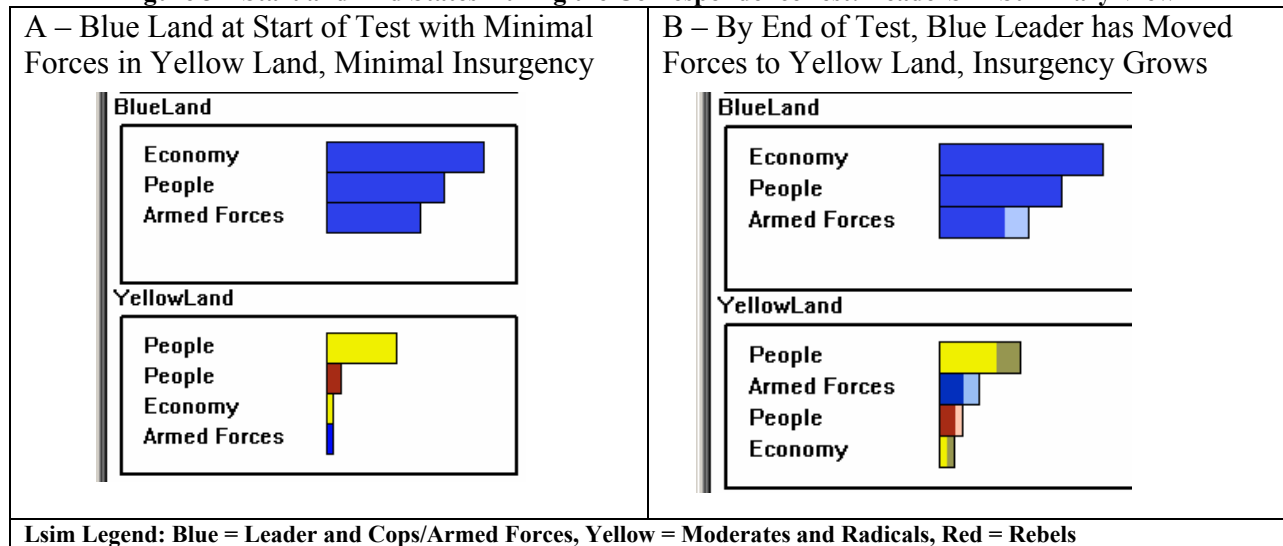


Figure 3 (cont'd) - Civil Violence View of Population Membership Before and After Correspondence Test

<p>C – Starting State (Avg of Weeks 1 & 2) Muslim Population at Start Is Neutral with Few Grievances Registering</p>	<p>D – End State (Avg of Weeks 103, 104) Muslim Population Reflects Radicalization and Spread of NonViolent and Violent Protest</p>
---	--

GrievanceState0 - Neutral	30%	6%	GrievanceState0 - Neutral
GrievanceState1 - Disagree	55%	1%	GrievanceState1 - Disagree
GrievanceState2 - Join Oppost	15%	37%	GrievanceState2 - Join Oppost
GrievanceState3 - Nonviolent	0%	39%	GrievanceState3 - Nonviolent
GrievanceState4 - Fight-Rebel	0%	17%	GrievanceState4 - Fight-Rebel
TOTAL	100%	100%	TOTAL

4) Socio-Cultural Game Results to Date: Turing, Correspondence, and Sensitivity 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. In this Section, we exam a scenario where a larger group of one religion and its leader (Blue) discriminate against 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? This case will draw training and test data from events during the past three years in SE 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 3A).




Figure 4: Grievance from Simulation and Observed Data (Incident Severity) [Training & Test Periods]

4.1) 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 3A 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, the Tsunami 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 4). The data was then longitudinally separated into ‘independent sets’ with training set consisting of Jan-June 2004 while test set beginning in July 2004 and running till Dec 2004. We curtail the test data to end before the tsunami.

Setting Up The Testbed 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.

In order to adequately test these PMFserv agents’ ability to interact at the population level, the PMFserv run groups are connected to a cellular automata that is known as the Civil Violence model (Epstein et al., 2001), though Leader Legitimacy is replaced with PMFserv agents’ view of membership. 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 or Yellow (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 writeup, the Red Group has no active agents, but is marked up as rebels that afford activations as mentioned in Section 3.

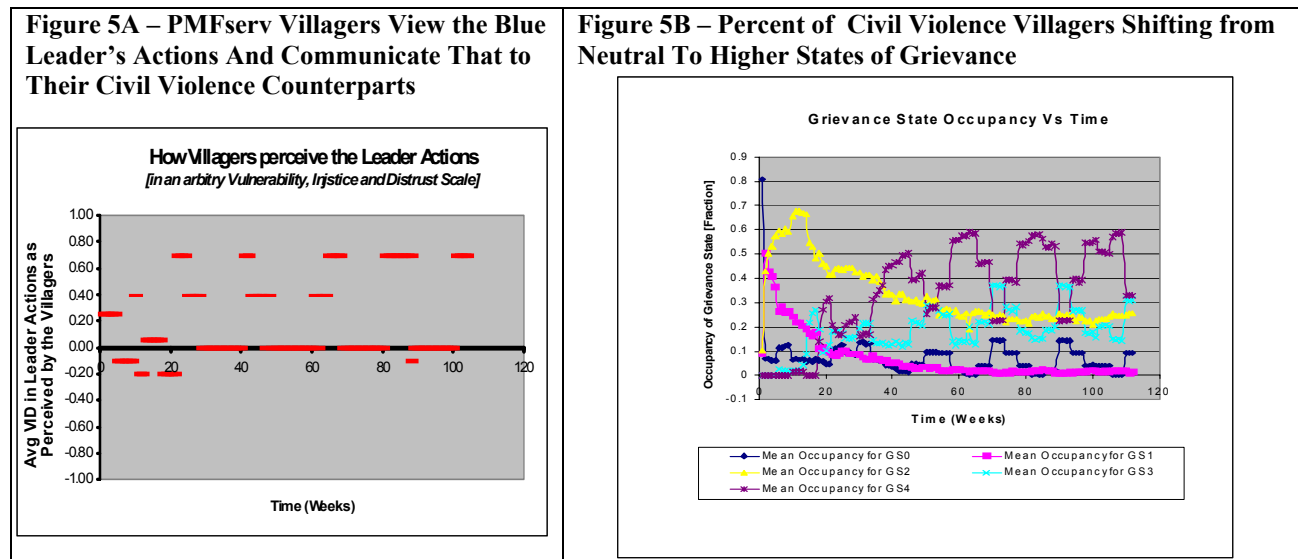
Running the Simulation -- The correspondence test is whether the overall parameterization for the GSP tree-guided PMFserv agents in the bridge with the Civil Violence population will faithfully mimic the test data set. That is, by tuning the GSP trees of 1 leader and 160 villagers, and by connecting 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.

The simulation starts on the left side of Figures 4A for Lsim and 9C 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 4B and D for each view, respectively. We can also examine what happened as the run proceeded. Figure 5A 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 5B 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. $IncidentSeverity = w_f \times fatalities + w_i \times injuries$. 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. 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



4.2) Validation

In order to assess the validity, we examine what transpires inside the heads of the various types of agents in the simulated world. 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 $H(\cdot)$ 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, indicating correlation between real and simulated data. Details of this validation as well as villager validation are given in the Appendix II. With an M metric, one cannot make statements about the confidence interval of the correlation, however, the Blue Leader in the current scenario seems faithful to his real world counterpart. This gives us reason to suspect the Hermann- and GLOBE-based GSP tree structure works equally well across time periods, locations, and cultures.

Not all aspects of the agent mindset work as well as this, however. As an example, one would expect agents to waiver somewhat in their resolve, yet our mechanism for that waivering seems like it may be too heavy-handed. The reader will recall that at the end of each cycle, the XML bridge feeds average Civil Violence neighborhood grievance back to the PMFserv agents in that neighborhood as a dampening of their reactions. This feedback is in the form of a replacement of the PMFserv agent’s memory of past grievance level. Resetting of past grievance leads to flip-flopping behavior of PMFserv agents. This causes moderate villagers to flip between supporting opposition (GS2) and joining government (GS-2) depending on what Blue Leader does, particularly during the early stages. Hence they seem overly fickle. Since radicals are more grieved, their flip-flops are between neutral (GS0) and fight back (GS4), and this seems less troublesome semantically, though a confirmed Jihadist probably has fewer of these issues. We are not sure if this amount of flip-flopping is warranted, and this could only be resolved by further behavioral studies of such individuals (studies of which may be underway by Atran (2006)). If research shows them to be less fickle in their inner beliefs, this is a relatively easy process to dampen in future versions.

4.3) Sensitivity Analysis

Once again, it is interesting to explore “what-ifs” and whether alternative decisions by Blue Leader will yield different outcomes amongst the YellowLand population. The reader will recall, however, that our model’s output is conflict parameters (action decisions, grievances, group membership), whereas the model’s inputs are characteristics of the leader and the followers. To change the outputs implies shifting the weights on the GSP trees of various archetypes of the population. Here we shift those weights for the Blue Leader, since we are interested to see if his personal decision style and choices are key to driving the villagers toward insurgency.

- By altering Blue Leader’s InGroup Bias we should be able to alter his decisions to provide more or less needed assistance (economic goods, non-Wahabbi schools), and then we can observe if that alters the outcome.

Specifically, we perturb InGroup Bias on his Standards Tree by 15% in either direction. Figure 6A shows the result.

- By altering Blue Leader’s Sensitivity to Life (Humanitarianism) we should be able to alter his decisions to provide more or less violent cops, and then we can observe if that alters the outcome. Specifically, we perturb SensitivityToLife on his Standards Tree by 15% in either direction. Figure 6B shows the result.
- By altering Blue Leader’s Openness we should be able to alter the immediacy of his response to opposition and protest. Thus he would send fewer cops down to YellowLand if he were more open, and more cops if he were less open, and then we can observe if that alters the outcome. Specifically, we perturb Openness on his Standards Tree by 15% in either direction. Figure 6C shows the result.

In examining these three sets of what-ifs in Figures 6A-C, as expected, a larger fraction of population occupies higher grievance states of 4 and 3, when the leader exhibits lesser degree of sensitivity-to-life and/or more InGroup Bias. Conversely, the population remains at lower grievance states when Blue Leader is more sensitive-to-life and less InGroup Biased. However, the trend is not the same with respect to the openness trait of the leader. It appears that more open leadership does not necessarily result in lower grievances in the community, but in less sustained (shorter) expressions of grievance.

Figure 6A – Alternative Outcomes (Grievance Level) When Altering Blue Leader InGroup Bias




Figure 6B – Alternative Outcomes (Grievance Level) When Altering Blue Leader’s SensitivityToLife

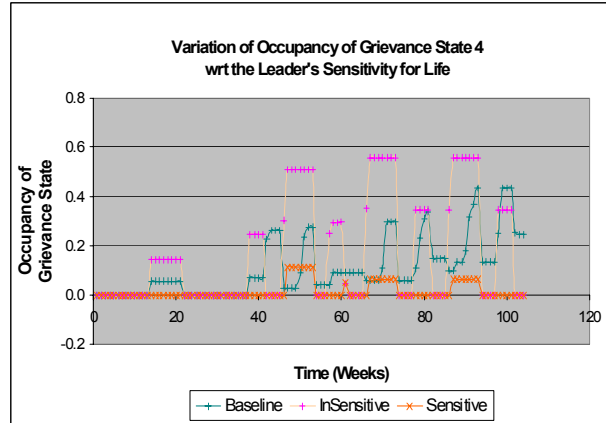


Figure 6C – Alternative Outcomes (Grievance Level) When Altering Blue Leader’s Openness





Figure 6D – The Low, Medium, and Upper Limits in Figs 6A-C Combine Into Points 1, 2, and 3 Culminating in a Tipping From Voice to Exit and Martial Law.

These types of results help us begin to calibrate the population’s demand curve mentioned earlier for exit, loyalty, and voice. Specifically, in Section 3 we presented several equations (equations 6-8) that help to determine the members’ decisions about expressing their grievance (voice) and/or exiting the legitimate authority and joining the separatist movement.

In Figure 6D, we see the graphical expression of several of those equations. We plot loss of civil rights and growing vulnerability, injustice, and distrust (VID) up the vertical, with group members along the horizontal. The demand curve is negatively sloped indicating people tend to increase their strength of membership in group A (BlueLand in this instance) as VID drops. The Transfer Rate helps to define the elasticity or slope of the demand curve. As the policies of the Blue Leader are altered for more Sensitivity to Life and Less InGroup Bias (Figures 6B and 6A, respectively), this is equivalent to shifting from point 2 on the demand curve toward point 1. On the other hand, more InGroup Bias and less Sensitivity to Life shift the YellowLand to point 3. In Figure 6D, we had the means to compute the desire to exit, and this is plotted along the horizontal axis. To the right of point 1, $\Delta\Phi$ is below zero, and few members favor the separatism. This is how YellowLand in fact was in the 1990s. In the interval between 1 and 3, the YellowLand people are indifferent since separatism cost is so high. There, they use their Voice and we observe the protests that occurred in the simulation runs of this paper, and that in fact occurred in the real world test dataset. Finally, to the left of point 3, Blue Leader's treatment becomes so intolerable, that he loses the hearts of the Yellow People and even the moderate followers are now on the side of the Red separatists or insurgents. In fact, in late summer of 2005, the Leader of this land had to declare Martial Law complete with curfews and movement checkpoints. The sign of our simulation results thus correspond well to the real world, and give us an ability to suggest outcome possibilities that are realistic for the Leader's policy choices.

5) Lessons Learned and Next Steps

In concluding, it is useful to revisit the three 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.

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. An example of its usage can be seen in SE Asian scenario described in this paper. 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, (2) 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.

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 alterative 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 (Aim 2). 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 Turning and Correspondence tests, where the leader and minority villagers 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 others often proved to be self-fulfilling prophecies. 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).

REFERENCES

1. Bharathy, G. K. (2006) *Agent based Human Behavior Modeling: A knowledge engineering based systems methodology for integrating of social science frameworks for modeling agents with cognition, personality & culture*. (Doctoral dissertation, University of Pennsylvania, July 2006)
2. Damasio, A.R. (1994). *Descartes' error: Emotion, reason, and the human brain*. New York: Avon.
3. De Marchi, S (2005). "Computational and Mathematical Modeling in the Social Sciences", Cambridge
4. Epstein, J., Steinbruner, JD, Parker, MT, (2001), "Modeling Civil Violence: An Agent-Based Computational Approach," *Proceedings of the National Academy of Sciences*. Washington DC: Brookings (CSED WP#20).
5. Hermann, M. G. (1999). *Assessing leadership style: A trait analysis*. Hilliard, OH: Social Sci. Automation, Inc.
6. Heuer, R. J., Jr. (1999). *Psychology of Intelligence Analysis*. Washington, DC: Center for the Study of Intelligence, Central Intelligence Agency.
7. Hirschman, A.O. (1970). *Exit, voice, and loyalty*. Cambridge, MA: Harvard University Press.
8. Hofstede, G., (2003), *Culture's Consequences, Comparing Values, Behaviors, Institutions, and Organizations Across Nations* Newbury Park, CA: Sage Publications. Also avail. at <http://www.geert-hofstede.com/>
9. House, RJ, Hanges, PJ, Javidan M, et al., (2004), *Culture, Leadership, and Organizations: The GLOBE Study of 62 Societies*, Thousand Oaks, CA: Sage Publications
10. Johns, M., & Silverman, B.G. (2001). How emotion and personality effect the utility of alternative decisions: A terrorist target selection case study. *Proceedings of the 10th CGF & BR, SISO*.
11. Johns, M., *Deception and Trust in Complex Semi-Competitive Environments* (Doctoral dissertation, University of Pennsylvania, 2006)
12. Ortony, A., Clore, G. L., & Collins, A. (1988). *The Cognitive Structure of Emotions*. Cambridge: Cambridge University Press.
13. Pew, R. W., & Mavor, A. S. (1998). *Modeling human and organizational behavior: Application to military simulation*. Washington, DC: National Academy Press.
14. Silverman, B.G., Johns, M., Weaver, R., O'Brien, K., & Silverman, R. (2002a). Human behavior models for game-theoretic agents. *Cognitive Science Quarterly*, 2(3/4), 273-301.
15. Silverman, B.G., Johns, M., O'Brien, K., Weaver, R., & Cornwell, J. (2002b). 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, Orlando, Florida*, 97-106.
16. Silverman, B.G. (2005). Human performance simulation. In J.Ness, D.Ritzer, & V.Tepe (Eds.), *The Science and Simulation of Human Performance (Chapter 9)*. New York: Elsevier.
17. Silverman, BG, Rees, R, et al. (2005, May), "Athena's Prism: A Diplomatic Strategy Role Playing Game for Generating Ideas and Exploring Alternatives," in *Proc. Internat'l Conf on Intelligence Analysis*, MacLean, VA: Mitre.
18. Silverman, B. G., Johns, M., Cornwell, J., & O'Brien, K. (2006a). Human Behavior Models for Agents in Simulators and Games: Part I – Enabling Science with PMFserv. *Presence*, v. 15: 2, April.
19. Silverman, B. G., Bharathy, G.K., O'Brien, K., Cornwell, J. (2006b). Human Behavior Models for Agents in Simulators and Games: Part II – Gamebot Engineering with PMFserv. *Presence*, v. 15: 2, April.
20. Silverman, BG, Bharathy, GK, Johns, M, Nye, B, Eidelson, R., Smith, T, "Socio-Cultural Games for Training and Analysis", Tech Report avail from the authors (submitted for publication 2006c).
21. Wood, EJ (2003). "Modeling Robust Settlements to Civil War: Indivisible Stakes and Distributional Compromises", Santa Fe, Unpublished Paper. Available from the author: wood@santafe.edu

APPENDIX I: Stylized Game Description

The agents in our model participate in a multi-stage, hierarchical, n-player game in which each class or type of agent (D_x) observes and interacts with some limited subset of y agents (human or artificial) via one or more communication modalities. We make three empirically plausible assumptions about multiple hierarchies of agents: (1) play multiple distinct games, (2) are cognitively detailed, (3) agents are self-serving, and attempt to maximize its utility (u) within this iteration of the game, as follows. **GAME = (a ∈ A, U_x, D_x)_ for x ∈ X** --[A1.01]

In this case, we set it up as a game (similar to that of Wood, 2003) between the Leader of the CountryT (Leader or L) and the minority villagers (V), who have conflicting-interests. The leader wants to divert his resources to his constituency, thereby discriminating against this minority. If they rebel, the leader would not hesitate to use violence (fight or f). If the leader were to compromise, he would not help his constituency in the short term (in turn may lose some support or votes especially from extremists), but would manage to settle the conflict (compromise or c). The outcomes are as follows for single shot and iterative games.

Single Shot:		Repeated (as negotiated solutions typically are): If future discount rate (rate of time preference) is i for both parties, then over infinite horizons:	
		Villager	
		Fight	Compromise
Leader	Fight	f_{1L}, f_{1V}	$f_{2L}, 0$
	Comp	$0, f_{2V}$	c_{3L}, c_{3V}
		Villager	
		Fight	Compromise
Leader	Fight	$f_{1L} \cdot (1+i)/i,$ $f_{1V} \cdot (1+i)/i$	$f_{2L} + f_{1L} \cdot [1/i],$ $f_{1V} \cdot [1/i]$
	Comp	$f_{1L} \cdot [1/i]$ $f_{2V} + f_{1V} \cdot [1/i]$	$c_{3L} \cdot (1+i)/i,$ $c_{3V} \cdot (1+i)/i$

Mutual conflict or fight-fight are Nash equilibriums, as compromising while the other player fights have a low pay off.
 $c_{3V} + c_{3L} = c_3$ and $c_{3L} = y \cdot c_3$ and $c_{3V} = (1-y) \cdot c_3$
 If $c_{3V} > f_{2V}$ and $c_{3L} > f_{2L}$, then mutual compromise is also a Nash equilibrium.
 This zero sum game assumes, fraction y of the total compromised divisible pay off (c_3) is allocated to the leader.

Initially negotiate or fight as per single shot game. Simultaneous, repeated game with information input from historic events. Mutual compromise or conflict, once established, tends to persists in the absence of exogenous shocks. For mutual compromise,
 $c_{3L} \cdot (1+i)/i > f_{2L} + f_{1L} \cdot [1/i]$ and $c_{3V} \cdot (1+i)/i > f_{2V} + f_{1V} \cdot [1/i]$

Let p_{LVC} and p_{VLC} be the probabilities estimated by the leader and villager for compromising by villagers and leader respectively. Using these, one may estimate the expected payoff for:

- the leader to compromise as: $E(LC) = p_{LVC} * c_{3L} \cdot (1+i)/i + (1-p_{LVC}) * \{f_{1L} \cdot [1/i]\}$
- the villagers to compromise as: $E(VC) = p_{VLC} * c_{3V} \cdot (1+i)/i + (1-p_{VLC}) * \{f_{1V} \cdot [1/i]\}$

Similarly, the estimated pay off for:

- the leader to fight is: $E(LF) = (1-p_{LVC}) * \{f_{1L} \cdot (1+i)/i\} + p_{LVC} * \{f_{2L} + f_{1L} \cdot [1/i]\}$
- the villagers to fight is: $E(VF) = (1-p_{VLC}) * \{f_{1V} \cdot (1+i)/i\} + p_{VLC} * \{f_{2V} + f_{1V} \cdot [1/i]\}$

Ranges of probabilities of belief required of opponent’s likelihood of compromise, over which a party will compromise, could be obtained as shown in the figure (robust equilibriums with symmetry). While such stylized games are useful to elegantly illustrate the division of tangible and divisible resources as a zero sum game, they do not take into account the intangible resources such as emotional and cultural pay off (as well as indivisible resources). On the other hand, our simulation models do handle these added factors.

APPENDIX II: Validation

This appendix provides a more detailed look at the runs and results between simulated vs. actual leaders and followers.

Correspondence between Simulated vs. Actual Leader Decisions

This correspondence test has been attempted with the use of the mutual entropy statistic. Our estimation of mutual entropy for CountryT is much less than 1, indicating that the real and simulated base cases might correspond well. Currently, there are no benchmarks that could indicate what would be an acceptable limit of mutual entropy for establishing correspondence. If we accepted an arbitrary limit of correspondence of mutual entropy less than or equal to 0.1 [an order of magnitude less than the mutual entropy associated with no correspondence], then we would be able to reject the null hypothesis of no-correspondence between real and simulated outputs. The following figure shows a summary of decisions of the simulated and actual leaders through the bar charts on the left and right respectively. The entropy calculations are shown below the charts.

Figure A2.1: Correlation of Simulated Leader vs. Real Action Decisions
Comparison of distributions to see Mutual Entropy (M). Reject H0 & Accept H1 if M<<0.1

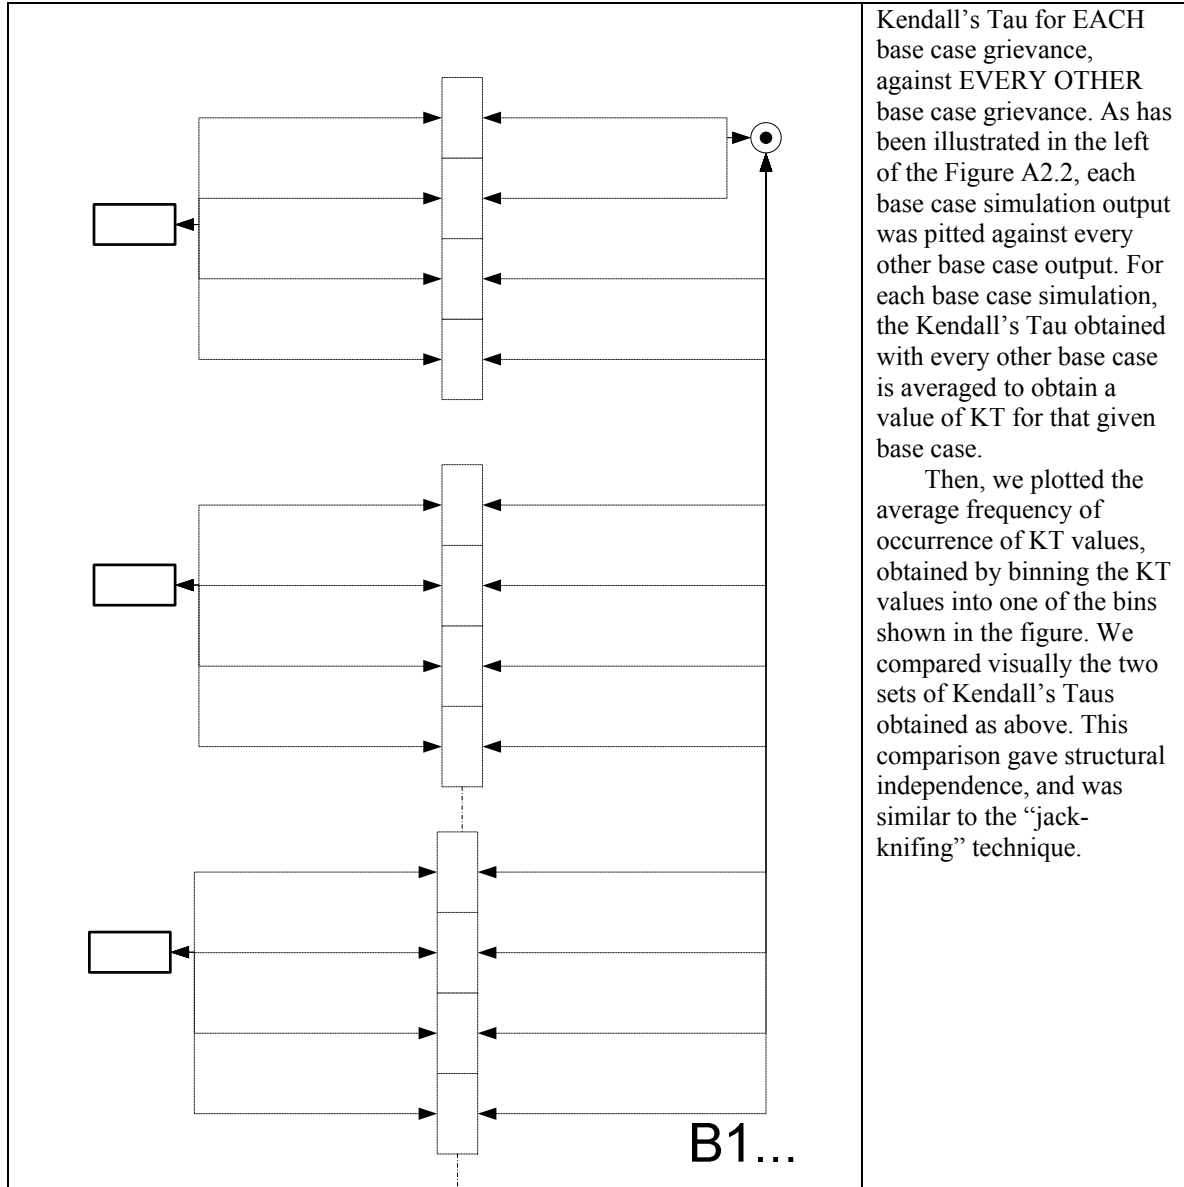
Distributions	PMFserv-Simulated Prime Minister's Actions		Real Leader's Chosen Actions	
Mutual Entropy Calculations	Joint Entropy of Sim & Real	1.396	$H(SIM, REAL) = - \sum p(sim_i, real_j) \log p(sim_i, real_j)$	
	Entropy of Sim	0.681	$H(SIM) = - \sum p(sim)_i \log p(sim)_i$	
	Entropy of Real	0.760	$H(REAL) = - \sum p(real)_j \log p(real)_j$	
	Mutual Entropy of Sim & Real	0.045	$M(SIM: REAL) = H(SIM) - H(SIM REAL)$	
Legend of Leader Actions				
Negative Actions: Discriminate Suppress - Increase Number of Cops Suppress - Increase Violence of Cops		Neutral Actions: Perceive (Observe Events)		Positive Actions: Give Culturally Sensitive Assistance Give Essential Assistance Reduce Suppress by Number Reduce Suppress by Violence


Villager Behavior Correspondence

We have employed Kendall's Tau to relate the real vs. simulated population's grievance data. Kendall's Tau is a measure of correlation, and so measures the strength of the relationship between two variables. It employs paired observations, and is scale-free. It is computed as the excess of concordant over discordant (nd) pairs, divided by a term representing the geometric mean between the number of pairs not tied on variable 1 and the number not tied on variable 2. There is no well-defined intuitive meaning for Tau -b, which is the surplus of concordant over discordant pairs as a percentage of concordant, discordant, and approximately one-half of tied pairs. We obtained two sets of Kendall's Tau (KT) values by comparing:

- simulated agent grievances from each base case simulation against every other base case simulation, and
- real population incident severities against the simulated grievances from every base case simulations.

Figure A2.2: Pairwise Comparison of Base Case Simulations and Real Data | We took averages of





It can be seen that KT values among base cases with moving averages are distributed in bins 3 (0.0-0.25), 4 (0.25-0.50), and - 6 (0.75-1.0), with proportions increasing in that order. This implies that a large majority of the pairwise comparisons (a little over half) among base cases result in high Kendall's Taus, while a smaller fractions have limited correlation.

Comparisons between real data (injury-fatality aggregated and smoothed, of course) and base cases (moving averages) are found in bins 3 (0.0-0.25) and 4 (0.25-0.50). This shows that all base cases show positive correlation with real data. A Kendall's Tau of 0.25 or 0.5 might appear to be a small correlation compared to a KT value of, say, 1.0, but in reality, these numbers indicate a fairly good degree of correlation, especially considering this is a time series and any mismatch would be counted as discordance resulting in negative correlation. This is illustrated by converting Kendall's Tau to p values, as seen below. Although the real data is an outlier, it is no more outlier than about half the simulation base cases themselves.

While we recognize that the p values are considered weak in the case of Kendall's Tau, and therefore exercise caution in the interpretations, it is hard to not notice p values for the same Kendall's Tau. The range of p values from the models run is less than 5%. If one were to take the approach of significance level, this would mean that one would be able to reject the hypothesis of mutual independence between the base case simulation grievance and observed incidents (villager or follower decisions) with a significance level of 0.05. With a two-sided test, we are considering the possibility of concordance or discordance (akin to positive or negative correlation). In our example, we can conclude with reservation (weak indicator, and hence the caution) that there is a statistically significant lack of independence between base case simulation and observed grievances rankings. That is, real incident data and simulation results are related. (From these range of results, and based on the alpha value of 5%, it seems reasonable to reject the null hypothesis).