Understanding *Why*–

Dissecting radical Islamist terrorism with agent-based simulation

Edward P. MacKerrow

The events of September 11, 2001, were a wake-up call for Americans. All of us remember where we were and what we were doing when we learned of the horrific attacks. On that terrible day, I was presenting results of a financial agent-based simulation to corporate executives on the top floor of a Houston skyscraper. I had no idea why the United States was attacked, but I had plenty of time to contemplate the question as I drove back to Santa Fe in my rental car. What had we done to motivate terrorists to take such hateful actions against us? What message were they trying to send us? Why did these attacks happen?

As the world strives to understand terrorism better and struggles to identify its many faces and forms, few would deny that religiously motivated terrorism is becoming increasingly prevalent. But the association of religion with violent, radical groups, many of which have their own interpretation of a religion, needs to be examined carefully. The violent actions of radical Islamist groups, for example, have led to the mistaken association of terrorism with Islam.

One-fifth of the world population is Islamic. Dispersed around the globe, the largest concentrations of Muslims are in Indonesia, Pakistan, India, Bangladesh, and the Middle East. In Western society, the largest concentration is in the United States. The religion itself is based upon peace (the word Islam means "self-surrender" in Arabic, and the universal greeting of Muslims is "*salaam alaikum*," which means "peace be upon you"). In many Islamic societies, however, the passive or neutral behaviors of the peaceful majority often become obscured by the attention-seeking acts of a "noisy minority." Although the point is debated, the general understanding is that most Muslims are peaceful because of their Islamic beliefs, and that the "noisy minority" has misinterpreted Islamic teachings.

Muslims are taught that Muhammad was sent by Allah to spread belief in a single God—as opposed to the multitude of pagan rituals honoring a variety of deities at his time (A.D. 610). For many modern-day Muslim radicals, especially those in traditional societies, American pop culture may be perceived as being similar to old-fashioned paganism, a cult that worships money and sex. Some modern-day militants may perceive themselves as following a path similar to Muhammad's in cleansing the Islamic world from the infiltration of the "pagan" West. Other Islamic people may fear that their culture, traditions, and beliefs are being replaced because globalization imposes Western values on them. Being mutually understanding of religious sensitivities, as well as responsible and respectful of each other's influence, will help establish a peaceful coexistence between the West and the Islamic world. Finding ways to move toward this goal requires careful analysis and discussion.

The Complex Systems Group at Los Alamos has been examining questions related to the "why" behind terrorist organizations in the Middle East. Borrowing tools from the field of computational economics and sociology, we are developing agent-based models that simulate social networks and the spread of social grievances within those networks. Our computer-generated "agents" are humanlike, endowed with personal attributes and allegiances that statistically match the demographics of a specified region and, like people, interact with one another and respond to societal pressures. The defining feature of our agents, however, is that their "behaviors" are allowed to change during a simulation run. For example, an agent may "learn" during the simulation not to interact with agents of a certain social class, or an agent may develop deep "feelings" of oppression and grievance based on its experiences.

We do not know *a priori* the life stories of our agents, but after tens of thousands have interacted, we have produced a scenario, or a virtual history, for a region of interest. The plausibility of this scenario is normally assessed by human experts who have complete knowledge of the model assumptions and the rules followed by the agents. By replaying any particular simulation, the experts can observe how the agents behaved and examine why they behaved in a certain way.

We can expose our agents to a variety of determinants—new government policies, different media exposure, economic pressures, and others—and quickly generate hundreds of new scenarios. Thus, we can conduct computational experiments that can be analyzed statistically and objectively to increase our insight, support decision making, and aid policymakers (see Figure 1). Scenarios can even be used to gain insight into actual events that have little or no historical precedence. It should be emphasized, however, that the goal of these simulations is not to predict *specific* events and not to estimate the probability or frequency of terrorist acts, but to generate scenarios and analyze them.

Our work is part of the Defense Threat Reduction Agency (DTRA) research effort known as the Threat Anticipation Program (TAP).[1] Built on a multidisciplinary team of scholars from the arts and literature, history and psychology, Middle Eastern and Muslim cultures, religion, economics, and sociology, TAP aims to develop algorithms and software frameworks that can generate the most likely models of terrorism and terrorist scenarios in order to catch the precursor signals of the next terrorist attack. We hope that TAP can eventually provide us with insight into potential prevention, interdiction, and mitigation policies.

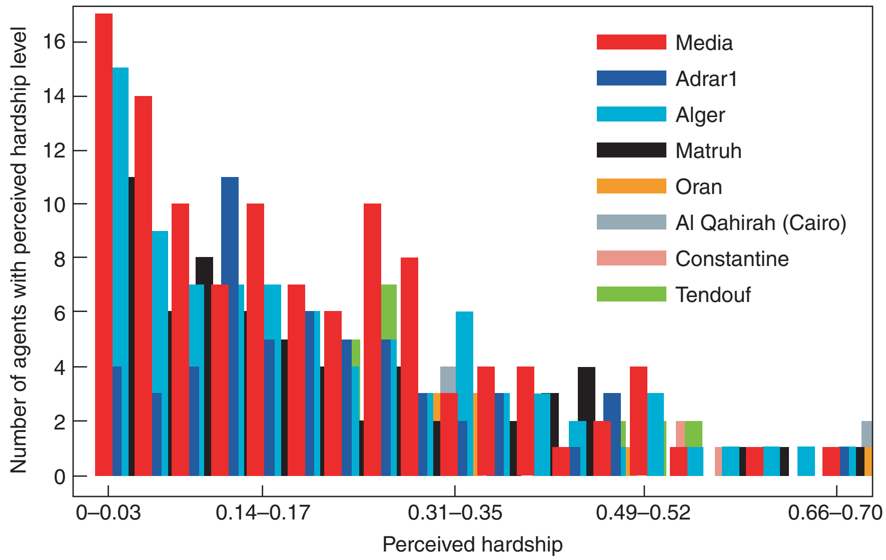

Figure 1. A Sample Distribution from the TAP Model

The TAP model places thousands of agents throughout the Middle East, endows them with numerous properties and behaviors, and allows them to interact for number of simulation years, thus creating a short, virtual history of the region. We can dissect that history and analyze various social metrics (such as perceived social disadvantage, directed grievances, and allegiances) to gain insight into extremist behaviors. The figure shows a distribution for perceived hardship, a metric that might be obtained in reality by a pollster asking the question, "On a scale of 0 to 1, what is your level of hardship?" The colors correspond to some districts in Algeria and Egypt. Because agent behaviors are not preprogrammed, each simulation using the TAP model will produce a different distribution of perceived hardship (or any other social metric). Human experts must assess whether the envelope of results is valid, that is, whether our model correctly brackets the possible levels of real-world social grievances. That task is difficult because the model can be validated with only one data point—the real-world poll made under specific conditions.

Modeling Complex Socioeconomic Systems

The social tensions in the Middle East emanate from many different yet interrelated conflicts, and each Middle Eastern nation has a unique history in relation to those conflicts (Miller 1997). Therefore, the underlying social processes cannot be understood by a simple linear combination of separate sociologic, economic, demographic, religious, cultural, and political subprocesses.

Agent-based simulation provides

[1]TAP was originated by Dr. Stephen Younger, director of DTRA. Younger is a former director of the nuclear weapons program at Los Alamos.

us with a methodology for modeling complex socioeconomic phenomena. Agent-based simulation was first introduced into economics to address shortcomings with economic simulations, which in early versions assumed homogeneous populations of idealized, perfectly rational agents who had perfect information about perfect markets. The results of those simulations, though frequently used, were often incorrect because of the flawed representation of real-world agent behavior. See Shubik (1997) for a discussion of issues associated with game theory applied to real-world applications.

As the computer became further integrated into the social sciences, more realistic socioeconomic models were attempted, and the methods of agent-based mathematics began to develop. Agent-based models have now found widespread use in economics and allow agents to act with bounded rationality, based on imperfect or incomplete information, and to act on chance and perceived economic utility. In addition to more realistic representations of individuals, agent-based simulation allows for analysis of nonequilibrium conditions compared with the historical practice of analysis made at equilibrium points—many real-world socioeconomic systems are not in equilibrium and may never reach equilibrium.

Currently, the major difficulty we face in building a model of a complex socioeconomic system is in quantifying social situations. Not only do we need better models that show how to represent social interactions, but we also need better empirical analysis of actual real-world studies. A fundamental problem is that real-world "observables" may be generated by many different interaction processes; therefore, empirical findings are open to different interpretations. My belief is that certain social micromodels apply better than others, depending on the context in which they are applied—and a good model will utilize a broad suite of social micromodels.

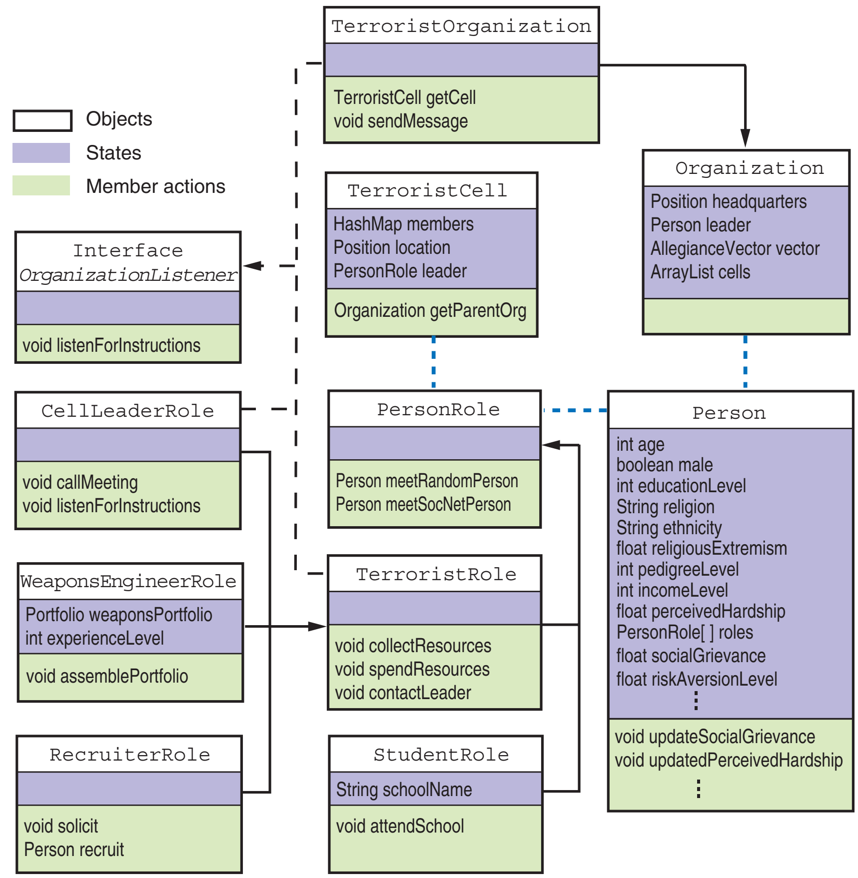

Figure 2. Objects in the TAP Model

Object-oriented programming is a natural fit to agent-based modeling because we can design objects that learn and adapt based on their history, their current state, and the states of other objects. This class diagram shows some of the object types used in the simulation. The simulation is built upon many different instances of these object types, each with different attributes. The object architecture allows for flexibility; the `PersonRole` class, and its inherited subclasses, allow a construct where any one `Person` object can play multiple roles. Interfaces allow for specification of required actions that can be implemented differently, depending upon the type of object implementing the interface (interface relationships are shown by black dotted lines). Objects can be composed of other objects allowing for new objects to build upon the structure of existing objects (shown by blue dotted lines).

The Los Alamos TAP Agent-Based Model

The TAP simulation is built with object-oriented software and implemented in the Java programming language. (The development of object-oriented software has been directly related to rapid advances in agent-based simulation over the last 15 years.) Although, in theory, any computer language can be used to represent a social system, object programming fits naturally with modeling social systems and greatly reduces development time.

In a simplified description, we first

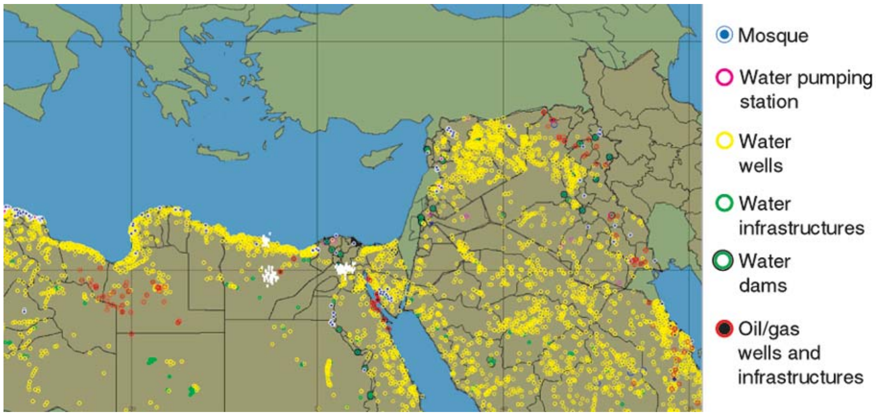

Figure 3. Region and Other Objects in the TAP Model
A screen capture from the TAP model shows many of the geographical region objects used to delineate different groups of agents as well as other objects. Each region is described by a different demographic composition. Agents are instantiated inside each region on the basis of these regional demographics. The white areas in northern Egypt, for example, indicate clusters of agents.

Table I. Attributes Used to Set the State of Agents

Agent Attribute	Distribution Type	Data Source
Age	Empirical (discrete)	U.S. Census Bureau
Sex	Empirical (male, female)	U.S. Census Bureau
Education	Empirical (years)	CIA World Factbook
Education type	Estimated (religious, secular)	CIA World Factbook
Ethnicity	Empirical (percent by group)	CIA World Factbook
Religion	Empirical (percent by group)	CIA World Factbook
Extremism	Estimated ([0,1])	Interviews and readings
Pedigree	Estimated ([0,1])	Interviews and readings
Income	Empirical ($ per year)	World Bank
Married	Empirical (Boolean)	The Economist
Employment	Empirical (Boolean)	The Economist, World Bank
Location	Empirical (# per km^2)	GIS Data Sources

establish the initial properties of "objects." The most important objects in the TAP model are `Nation`, `District`, `Organization`, `Mosque`, and `Person`. (In this article, we identify all software objects by capitalization and by using a Courier font.) Each object contains attributes that support a corresponding abstraction representation and modes of actions that represent agent behaviors (see Figure 2). We divide each Middle Eastern nation into a series of administrative `District` objects, as defined by standardized Geographical Information System data (see Figure 3). Each `District` is then populated with a number of agents. Empirical distributions derived from regional demographic and ethnographic data (see Table I) are used to initialize agent attributes. Different types of agents are instantiated with different data sets and different rules of behavior.

For each `District`, we define relative weighting factors to estimate the "social welfare," or "social capital," of religious and ethnic groups in that region. We posit that social capital is a weighted sum of income, ethnicity, religion, education, and pedigree, where pedigree represents inherited or appointed social wealth. (A Saudi prince would have a very high pedigree value.) We allow for flexibility of social weighting factors across different districts; for example, in some districts, ethnicity may not influence social status as much as does income.

We also take into account the important aspect of social rank in a society. Each agent updates both social capital and social rank on a regular basis during a simulation run. As discussed later in the text, social rank plays an important role in the theory we use to model social interactions—although by itself it is not necessarily a determining attribute of a terrorist.

Modeling Interactions

Many of the behaviors and opinions of individuals are rooted in the social structures to which they belong. For example, a young adult's proclaimed dislike toward the United States may be socially inherited from parents and reinforced through family circles, friendship circles, and by the media. Understanding which social experiences, conditions, and interactions increase the likelihood of becoming a terrorist is very difficult. Lay explanations for an individual's choice to become a terrorist include desperation, poverty, mental illness, and lack of education. Statistics of suicide attackers, however, show that low income levels are neither necessary nor sufficient to explain suicide attacks. Education is not a determining factor either, and ironically some data may suggest that higher education levels are positively correlated with terrorist attributes. So what is it then? The answer may be found by a closer examination of the terrorist organizations themselves.

Many terrorist organizations use charismatic leaders to cultivate and indoctrinate small cells of young recruits to become martyrs for the overall cause (Atran 2003). (Note that, for terrorist organizations, the use of martyrs is very economical, approximately $150 per suicide bomber attack.) These leaders, who rarely become martyrs themselves, seek to establish a commitment between the members of the cells in the form of a social contract—usually sealed with a video testimony. A cell member develops a sense of obligation to the fictive "kinship" of the cell. The social cost to individuals for reneging on commitments is very high; they risk being labeled "kafir"—an infidel, or nonbeliever. Defectors from some terrorist organizations can even be killed by the organization. Peer pressure from the cell is a motivating factor for some terrorists; others are fully committed on their own and believe strongly in becoming terrorists. The latter appears to be particularly true with respect to terrorists who oppose Israeli settlements in the occupied territories of Palestine. In the TAP model, we include different types of agents and allow some agents to be more self-motivated toward terrorism and others to require more or less active recruitment.

Structural and institutional constraints induce individuals to act in a manner most consistent with the preferences of the social structure or institution—in this case, the terrorist organization. In the TAP model, we introduce these constraints using social network representations for establishing probabilistic social rulesets. To include realistic interactions, we construct several different types of social networks between agents: kinship, religious, organizational, and friendship. Although meetings can occur between any two agents or between an agent and a group, we assume that interactions are more probable between agents on the same social network than between random agents in the local population, and we weight the interactions accordingly.

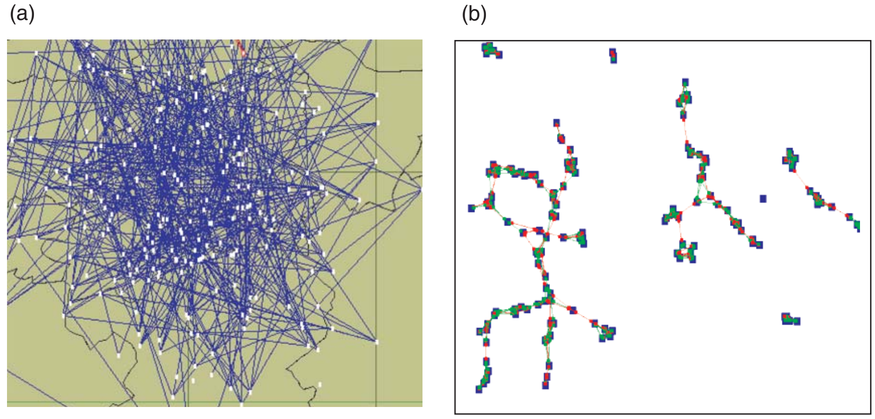

Figure 4. Evolving Networks
(a) This screen capture shows a large friendship network of agents in the TAP model at the end of the preprocessing stage. At this point in time, the network is rather dense; each agent is connected to an average of five other agents. A few agents initialized as Armed Islamic Group (GIA) terrorists are shown in red. (b) As agents interact with other agents during the simulation, the network evolves. The colors of the lines connecting the agents reveal how the agents became friends—either through random meetings or through mutual friendships. Some agents become isolated in stranded cliques and associate with very few other agents.

A screen-capture image of an initial friendship network of agents is shown in Figure 4(a), where we have used the algorithm of Jin et al. (2001) to construct the network. It is obvious how dense the friendship network is at the end of the preprocessing stage—each agent has a large number of social contacts. During the following stages of the simulation, this network will evolve into cliques and a less uniform density as agents evolve their social networks—as shown in Figure 4(b).

We include simulated social networks as structures that affect social interactions in our model, although we have also included the ability to populate the model with known terrorist networks. Obtaining data and characteristics on the actual social structures of the various terrorist organizations around the world is difficult because of their covert nature. In place of existing terrorist networks, we use surrogate information to model these social networks. (An example of a surrogate network is given in the box, "An al-Qaida Network" on the opposite page.)

Interacting Agents

When an agent meets another agent in our simulation, a social interchange occurs. Depending on the outcome of the meeting, a certain amount of interactive learning occurs. Because interactions are part of the larger social structure, a form of "social learning" occurs throughout the agent population—observable in part through their heterogeneous allegiances.

Each agent in the TAP simulation carries an "allegiance vector." The elements of this vector contain an integer value representing that agent's allegiance toward specific nations, organizations, ethnic groups, or religious groups. A positive (or negative) value of allegiance suggests a positive (or negative) allegiance for the group associated with that element. Figure 5(a) shows a schematic of the alle-

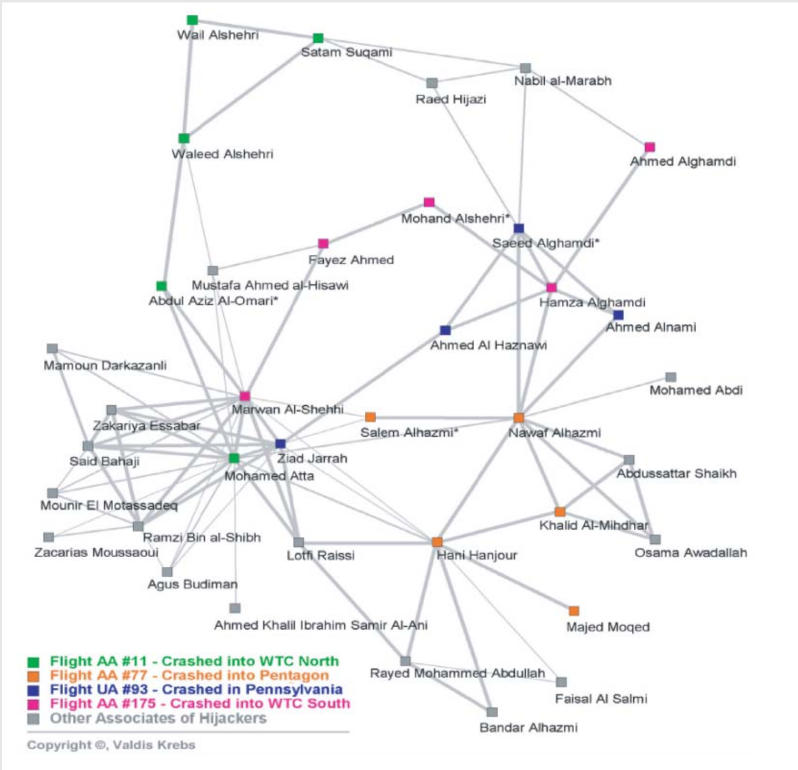

An al-Qaida Network

The compartmentalized social network of the al-Qaida cells involved in the September 11 hijackings, researched by Valdis Krebs (2002), is shown above. Krebs obtained open-source data on the hijackers as those became available after the attacks. Although some nodes are likely missing, his analysis and construction of the network are very useful in understanding some features of terrorist cells.

The social network of the hijackers was very loosely connected and sparse. Whereas Mohamed Atta is clearly seen as the ringleader, many of the hijackers were separated by a few degrees—more than one step away from each other. This separation even applied to hijackers on the same flight. The strategy ensures that the entire network is robust to the capture or compromise of a cell member. Usama bin Laden described this strategy on a videotape that was found in a deserted al-Qaida house in Afghanistan. This type of covert social network suffers from reduced informational efficiency and information sharing, although the hijackers were clearly able to mitigate those deficiencies and ensure some level of communication and resource planning. (Courtesy and permission of Valdis Krebs.)

giance vector concept.

The Gallup Polls of the Middle East are aids to understanding the way Islamic nations feel toward other nations on a variety of topics. We use these polls in the TAP model for estimating our aggregate allegiance vectors. Figure 5(b) shows an example of data from a Gallup poll.

Social Bargaining and the Nash Demand Game. Many theories of social interaction can be used for modeling ways in which beliefs, opinions, and values are communicated, shared, and modified during social meetings. We require a theory that allows us to model how allegiance values are transferred between interacting agents.

One theory we use is a social bargaining theory, established by H. Peyton Young (1998), that is based on the one-shot Nash demand game. During an interaction, each one of two agents places a bid for some portion of an abstract available "property," where the bid is related to the estimated value of establishing a social contact with the other agent. Both parties get their demands if the sum of the bids is less than the total property available; that is, successful bargaining occurs if the two agents are not too greedy. If one agent's bid is b_i and the other's is b_j, then a successful bargaining process occurs if $b_i + b_j \leq 1$. The condition for a pure Nash equilibrium (where the agents are at their best-bid positions, and a change in bid by either agent will lower the overall payoff) is $b_i + b_j = 1$.

Agents learn from past interactions how to bid optimally in Nash demand games. Each agent retains a memory of its past *m* meetings. Agents associate some attributes of those agents that they interacted with in these past *m* meetings and judge how well they did in demand games with agents of similar attribute types. This information is then used as a basis for future bids with agents of similar type—in a form of social learning.

If the one-shot social bargaining game is successful, then allegiance values are transferred in an asymmetric, bidirectional manner between the two agents. The agent with the lower social status "absorbs" more of the other agent's allegiance values, whereas the agent with the higher status absorbs fewer. In this way, social norms can emerge from the continued interactions between agents. Note that this is but one heuristic from a set of “social rules” used by the agents.

We hope that this methodology will represent one component related to the spread of beliefs through social structures. Some Islamic organizations establish relief efforts and supply resources for citizens of low socioeconomic status. In addition to providing needed welfare, these actions also improve public support of and opinion about Islamic organizations. Some organizations rely upon a “bottom-up” approach to instill the populace with their doctrine (that is, Islamic Law—Shari’a), whereas others may choose a more revolutionary “top-down” approach by replacing a secular regime with a more Islamic ruling party. The social bargaining approach fits closer with the “bottoms-up” approach of belief spreading.

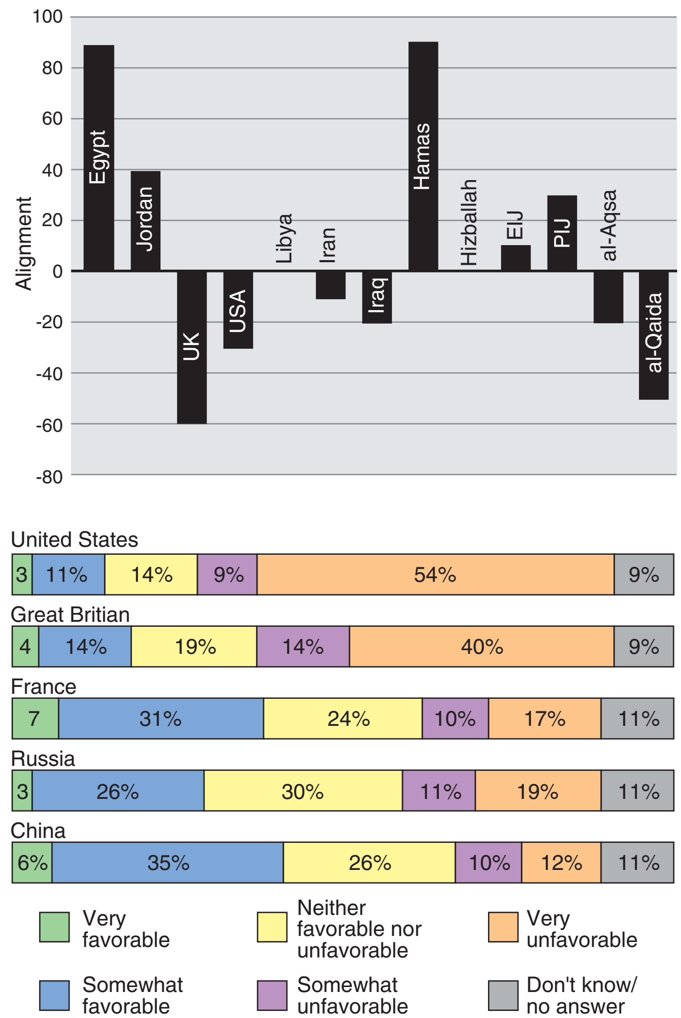

Figure 5. Typical Allegiance Vector
(a) All agents have allegiance vectors, which contain more than 100 elements. The values of individual vector elements indicate the feelings of an agent with regard to countries, organizations, and groups. These values change as agents interact with each other, thus simulating the process of social memes or “contagion” effects. (b) The results of Gallup Middle Eastern opinion polls help us quantify allegiance values in our model.

Social Repression and Estimation of Social Grievance

I have briefly outlined how allegiances are transferred through agent populations by social learning. This learning is affected by who meets whom, how socioeconomic status is valued relative to allegiance adoption that occurs between agents, and past experiences of agents in meeting with agents that are considered similar. At the end of the simulation, however, we want to quantify relative measures of propensity toward terrorism across the agent population.

Social grievance is one of the “summary” metrics we use to determine when an agent, or a collection of agents in a region or organization, is considered to have a propensity for protest and therefore has a higher potential to become a terrorist. It is calculated from metrics that include an agent’s sense of social repression. Social grievance is directed against a particular group or organization. In the simulation, repression stems from social disadvantage, inherited allegiances, cultural penetration, repression from the regime, and media influences.

We calculate a composite socioeconomic disadvantage metric based on the agent’s social ranking and the social ranking of the groups that the agent identifies with most. An agent may have high social status but may identify with a social group that has low social status—thereby increasing the level of disadvantage felt for other members of the group.

Agents perceive repression by a corrupt regime. We quantify this level of repression by calculating the overlap between the agent’s and the regime’s allegiance vectors, weighted by a corruption factor. Corrupt governing regimes in the Middle East that are secular and aligned with Western nations for commercial trade and military support, and who do not support the underlying social welfare of their populations, are considered to be an important root cause of militant Islamist terrorism against the West.

The infiltration of some aspects of Western culture, whether physically or media based, into Islamic regions is considered to be another source of repression to Islamic groups. In the TAP model, we estimate the contribution of cultural penetration to social repression as a weighted sum of the fraction of the local population that is nonindigenous and a media presence factor, which represents the relative amount of external (foreign) media influence in that region. The weighting function depends on the relative allegiance values—influences from cultures with high allegiance values are considered good. We estimate the media influence factor from "surrogate" data. For many Middle Eastern nations, we have measures of the percentage of households that have access to television, radio, and the Internet. Though not always the case, a household that has access to these information sources usually has access to the media of all other cultures. (In some Middle Eastern regions, the regime may censor external media influences.)

After calculating an agent's social disadvantage and level of repression, we calculate agent A's time-dependent grievance G toward the social group M as

$$G = d(t) \times a(M, t) \times o(t) \times p(t) \times f(M),$$

where the first term is the socioeconomic disadvantage of A, the second term is the dislike A feels for group M, the third and fourth terms are the regime and cultural-penetration contributions to A's perceived oppression, and the last term is a measure of A's perception of group M's level of corruption. Social grievance of the agents in each region is monitored in order to obtain a social grievance potential as the simulation progresses. An indicator of terrorist instability exists if agents with high levels of social grievance have access (through their social networks) to a terrorist organization with similar grievance targets. This represents increased public support and more probable recruitment for terrorist organizations.

Summary

I have briefly described a complex agent-based model of a complex situation—terrorism associated with militant groups. My hope is to give a flavor of the methodology used in constructing agent-based simulations of socioeconomic systems and to show how this methodology is being applied to the challenges in developing a detailed understanding of the sociodynamics of militant Islamist terrorism.

I see agent-based simulations as computational experiments that convey a great deal of scenario information in a timely, efficient, and safe manner. Examining the path-dependent time evolution of a particularly interesting simulation result (a virtual history) in replay mode allows us to analyze the "how" and "why" of agent behavior. If the TAP agent-based model can supplement policymaking in the turbulent clash between the West and Islamist radicals and can help policy makers visualize and understand important yet currently unknown interrelationships, then this work will be a success. Although agent-based simulation can help us gain insight into complex system behavior, my longer-term hope is that the "tit-for-tat" pattern of violence in the Middle East will be quenched by taking an honest look at both sides of the issues at hand through peaceful negotiation and mutual respect.

Acknowledgments

I thank the DTRA Advanced Systems and Concepts Office for their kind support and valuable expertise in constructing the TAP model. I also thank Gary Ackerman, Anjali Bhattacharjee, and many others at the Monterey Institute of International Studies for sharing their valuable expertise on terrorism and the Middle East; Charles Macal and Michael North of Argonne National Laboratory for their expert modeling advice; and the Santa Fe Institute for its support and useful input; Amir Mohagheghi, Faraj Ghanbari, Jessica Turnley, and Nancy Kay Hayden of Sandia National Laboratories for their useful discussions. I thank Waleed El-Ansary of the Islamic Research Institute for his useful insights into Islamic economics and game-theory. Finally, I thank the Los Alamos team of Joe Holland, Karl Lautenschlager, Merle Lefkoff, Ayla Matanock, Dennis Powell, Brian Reardon, David Sharp, and Zoltan Toroczkai for their important contributions to this work. ■

For further information, contact Edward MacKerrow (505) 665-3491 (mackerrow@lanl.gov).